

hopper (Mark Hofmann)
Friday, June 30, 1995 7:32 PM
tbr (Tim B. Robinson)
wampler (Kurt Wampler)
sisyphus tape out

Hi Tim,

We notice that we omitted to get rid of the MWAP layer over the cache area on the calliope baseplate. The result is that contped and metall did not get waffled in that area. Al is concerned that this large hole may peel and harm other active area. Therefore we are going to start a re-fracture of these 2 layers tonight.

The results will be ready for shipment Monday (and already have Al and Paul's blessing to just ship out the tapes.) will you be available to review the paperwork on Monday? Kurt and I discussed sending an postscript file of the paperwork to you for your review if you were planning on working remotely that day.

I haven't been able to talk to Grant (as a back up)- he's been in a meeting. So I'm not sure if he'll be around Monday.

-thanks,
hopper

From: geert (Geert Rosseel)
Sent: Friday, June 30, 1995 1:18 PM
To: pollux
Subject: Pollux status

Hi,

Here is the plan for Pollux :

1. Check-in a couple of new layouts (Tom)
2. Start building the Euterpe/Proteus snapshot today (Tim)
3. Release the Pollux toplevel BOM by tomorrow (Geert)
4. Tar the Euterpe/Proteus snapshot and
build a Pollux snapshot (Tom)
5. GETBOM the Pollux snapshot (Lisa)
6. Build the Pollux snapshot (Geert)
7. Start-up a toplevel LVS/DRC (Geert)

These tasks are all dependent on each other. We'll have to page/e-mail each other when a task is done. We should get it done by the end of the long weekend (Tuesday evening).

There are a couple of open issues :

- > crack-protection ring. We need this before we can tape-out.
Paul : Can you make this your highest priority ?
I need it today.
- > logic verification on mod8 not clean
- > logic verification on mod10 not clean
- > Most cells still have DRC errors.
This should not stop us from building Pollux. We'll
build Pollux using the current layouts and then
decide on what changes we want to pick up in the
snapshot.

Geert

From: doi (Derek Iverson)
Sent: Thursday, June 29, 1995 8:11 PM
To: guarino; gmo; iimura; claseman; jeffm; ethan; lisa
Cc: hestia
Subject: Software Bringup Meeting Minutes - June 28, 1995

Software Bringup Meeting

June 28, 1995

Next Meeting: July 5 at 2:00 pm.

Note: Loretta and Tim will be absent from this meeting.

Attendees: guarino, gmo, doi, iimura, claseman, jeffm, ethan, lisa

New Action Items

Item: Run OSF kernel on the IKOS

Who: lisar

Status: New

06/28 There are a number of verification tests that need to be run
on the IKOS before we run the OSF kernel.

Item: Enable SW folks to use the HW simulators

Who: doi

Status: New

06/28 doi is to start working on a queing system to enable others use the
HW simulators.

Item: Add support to terp to have the proper cerberus power-on defaults.

Who: Ethan

Status: New

Item: Analyze why regdepend tests take longer on the HW than terp thinks.

Who: dit00, claseman

Status: New

Item: Modify the ukernel's taskExit so we can tell when they finish on
the HW simulators.

Who: gmo

Status: New

Review of Action Items

Item: Plan for testing remote debugging environment

Who: everyone

Status: Pending

04/13 This will be discussed at the next meeting.

04/20 Short discussion on the pieces required and if we want to have
a standalone version of the hostio software.

Wally has started on this already.

04/26 Any work with snoopy/dram models will begin after euterpe tapeout when jeff has more time.

05/03 Wayne F. expects to have a prototype ISA card ready for testing tomorrow. Gmo is in the process of writing code to test the card.

05/10 Gmo has written a test program. Gmo and Wayne need to get together and try it on the board.

05/24 Wally has added more connectivity between the ukernel and the debugger (ability to interrupt kernel via the cerberus 'forced interrupt' bit).
Gmo is ready with a test when Wayne has the board ready.

05/31 Wayne has card ready to test (fpga that talks to ISA) for initial testing. Documentation for the interface next week some time.

06/14 Gmo has been able to access Wayne's board over the ISA bus. Informal review of the ISA CBI device later today.

06/28 CBI document has had initial review. Wally is working on getting gdb to run using the umbilical cord.

Item: ukernel needs to detect machine checks and 'do the right thing'
Who: gmo
Status: Pending

04/13 No new progress.
04/20 On list of things to do. Lower priority.

Suspended Items

Item: Unsnap code
Who: lisar, jeffm
Status: Suspended

02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned.
Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump.
We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.

03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.

05/10 Back to life. Does the IKOS support RAM dump?

05/24 This item was suspended again. There are no resources allocated to this item at this time.

Item: When do we have a full calliope simulation available (IKOS)?
Who: lisar
Status: Suspended.

04/26 This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.

05/10 Lisar was not able to attend the meeting.

05/24 Suspended. This is related to the snap/unsnap item (above).

Completed Items

Item: Build osf kernel and perpare it for execution on HW simulator

Who: doi/iimura
Status: Done.

06/14 Derek is going to send e-mail to lisar on how to accomplish this.

Item: Create performance test plan

Who: claseman

Status: Done.

11/30 We continue to run tests to help us compare terp vs hardware performance.

We still need to put together the actual performance tests that need to be run on the hardware.

05/10 Back to life. Tim has checked in more tests. We are going to use the tests written for hermes and cerberus read accesses in the 'How to debug Euterpe' presentation on monday. We will be looking for the time it takes from the instruction issue, entry to nb, and completion of access.

05/24 Tim is going to send lisar a list of the test names that he intends to write so they can be incorporated into the test template.

05/31 Jeffm is going to analyze approx one test per day (with Tim looking over his shoulder) to determine/verify the actual performance numbers so they can be incorporated into the software simulator.

06/21 Jeff has published numbers for dram load and a dcache fill.

Terp Feature Status

inprog o Ifetch protection granularity
 - performance vrs accuracy tradeoff
inprog o Fetch instructions as octlets
 o Implement hardware configuration through Cerberus regs
 (SDRAM parameters, dram enable?)
done o remove stbio from hwterp.
new o power-on cerberus defaults

----- line of conciousness -----

o Better latency model for Calliope accesses
o Checkpoints/Snapshots

Performance Test Status

Status Test Description Verified H/W S/W -----
ran store/load to dram Yes No
ran store/load to hermes No No
ran store/load to cerberus Yes No
ran load from rom Yes No
ran icachemiss No No
ran dcachemiss Yes No
 load ltlb entry (write+read) No No
 load gtlb entry (write+read) No No
 NB overflow No No
 generate an interrupt No No
 return from interrupt No No
 mult cyls trying to take exceptions at the same time No No
 predicted branch No No
 unpredicted branch No No
 single cylinder exception No No

ran cable_in_main_handler-- inner loop in No No

EQ_UPDATE_WEIGHTS() (khp)			
NTSC encode loop (ron)	No	No	
rs_compute_syndrome() (larry)	No	No	
IDCT code (gregg)			
pseudo-Huffman decode loop (gregg)	No	No	
a reconstruction routine from macroblock.c (gregg)		No	No

Classes	Tests Written
-----	-----
Loads	
Stores	s64la
Atomic	saas64la, smas64la
Branches	bi
Gateway (??)	bgate, bgatei
blink{,I} (pred and unpred)	--
{G,E}set	gsete64
{G,E}{add,sub,mux,logical}	eaddi
Easum	easum
Eulms	--
Elms	--
Ggfmul8	ggfmul
G{,U}mul{,add}{8,16}	gmuladd8, gmuladd16
G{,U}mul{,add}32	gmuladd32
G{,U}mul64	--
G{,U}muladd64	gmuladd64
Gextgractl128	--
Other gextract instr	gextractil128
Expand/Compress	--
Shift/Rotate	--
Copy/Swap	--
Shuffle/Deal	--
Suffle/Mux	eshufflei4mux
Select	--
Deposit/Withdraw	--

Dave Tomlinson has noticed that the stgen (store/load) tests have similar cycle counts between the HW and SW simulators but the regdepend (register dependency) tests take 2 - 3 times as many cycles on the HW as predicted by the SW simulator. Jeff and dit00 are going to analyze a trace.

Test Status and General Discussion

There are two outstanding event-daemon problems.

- o duplicate blocking reads are being allowed
- o duplicate caused the wrwong NB entry to be cleared

Verification is going to re-run all the tests to get up-to-date results.

Loretta was wondering if we are going to use the TCC compiled versions of the acceptance tests.

Ethan would like to know each time a test does not run correctly (or someone believes a test should not run) on the software simulator. This way he can either fix the problem or verify that the test will not run if it requires support that terp does not have (like some tests that rely on a special 'test mode' that is enabled in some behavioural models).

Amber L. Stone

From: lisar (Lisa Robinson)
Sent: Wednesday, June 28, 1995 10:43 PM
To: tbr; mws; dickson; dit00; jeffm; billz; woody
Cc: euterpe
Subject: Randomly generated register dependancy tests

Due to my misunderstanding of how the randomly generated register dependancy tests were implemented these tests will all require re-run and previous (Ran ok) status of these tests should ignored. The status file will be purged of these tests. (The test for ran-okayness was incomplete.)

These tests (regdepend_r???) do not feature in list of "required for tapeout" nevertheless their re-run will be completed as a high priority.

Note that this does not effect the status of the randomly generated store/load/execute tests (stgen_r???) .

Lisa R.

From: hopper (Mark Hofmann)
Sent: Wednesday, June 28, 1995 1:05 AM
To: Tim B. Robinson
Cc: dickson (Richard Dickson)
Subject: Re: Help! dead in the water

Tim B. Robinson writes:

My top level euterpe run is dying with pim2pif errors. The problem seems to be spacer cells overlapping. However, what's more ominous is the pim file is full of floating point numbers, which have never been there before. For example:

```
xlu/g_ctrlldata/g_d_6ax_0      128      2      xlu/input
xlu/g_ctrlldata/g_xbus2_7ax_0    128      3      xlu/xbus
xlu/g_ctrlldata/g_xbus0_7ax_0    128      4
0s 128      0.969397
0s 128      0.969397
0s 128      0.969397
0s 128      0.969397
0s 128      0.969397
0s 128.001    0.969397
0s 128.001    0.969397
0s 128.001    0.969397
0s 128.001    0.969397
0s 128.001    0.969397
0s 128.001    0.969397
```

[snip]

Okay. The floating point numbers are not an error by themselves. They are introduced by the AWK scripts when the row and column numbers are re-written.

It's normal. It's supported by the PIM/PIF syntax.

The duplicate cells are errors, as you found out.

As far as I know these should not occur with the usual make flow. If a top-level .pim file was assembled by hand from several lower level .pim files then this could happen. Do you know how the top-level file was created?

```
.xoffset 1242
.yoffset 35
0s 0.9995      8.06981
0s 0.9996      8.06981
0s 0.9997      8.06981
0s 0.9998      8.06981
0s 0.9999      8.06981
0s 1          0.99981
xlu/g_ctrlldata/g_cla_5ax_21      1      1
```

The xlu sub-block appeared to run OK by itself, but the pim file there is also full of 0s spacers and floating point numbers. The problem, I think is a rounding/truncation problem somewhere when the top level file is being constructed. However, what are all these, and why have they suddenly appeared.

This part of the file looks fine. Tom Vo has an intricate xlu layout which essentially combines a section of fingers and spaces with another section of finger and spaces- the fingers of one section fit into the spaces of the other section. He uses lots of spacers. I think these spacers have been at top level before- you just haven't noticed them! They do change from iteration to iteration.

Tim

-hopper

From: tbr
Sent: Wednesday, June 28, 1995 1:46 AM
To: hopper
Cc: dickson
Subject: Help! dead in the water

My top level euterpe run is dying with pim2pif errors. The problem seems to be spacer cells overlapping. However, what's more ominous is the pim file is full of floating point numbers, which have never been there before. For example:

```
xlu/g_ctrldata/g_d_6ax_0      128  2      xlu/input
xlu/g_ctrldata/g_xbus2_7ax_0  128  3      xlu/xbus
xlu/g_ctrldata/g_xbus0_7ax_0  128  4
0s      128      0.969397
0s      128      0.969397
0s      128      0.969397
0s      128      0.969397
0s      128      0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.001    0.969397
0s      128.002    0.969397
0s      128.002    0.969397
0s      128.002    0.969397
0s      128.002    0.969397
0s      128.002    0.969397
0s      128.002    0.969397
.xoffset 1242
.yoffset 35
0s      0.9995      8.06981
0s      0.9996      8.06981
0s      0.9997      8.06981
0s      0.9998      8.06981
0s      0.9999      8.06981
0s      1          0.99981
xlu/g_ctrldata/g_c1a_5ax_21    1      1
```

The xlu sub-block appeared to run OK by itself, but the pim file there is also full of 0s spacers and floating point numbers. The problem, I think is a rounding/truncation problem somewhere when the top level file is being constructed. However, what are all these, and why have they suddenly appeared.

Tim

Amber L. Allen:

From: tbr
Sent: Wednesday, June 28, 1995 12:37 AM
To: gmo
Subject: OSF question

We are seriously considering a big DEC alpha machine as a possible way of speeding up the tapeout fracture critical path. Based on the data we have so far it looks like it could be up to 4x the current trex machine in integer performance (we don't care about fp). It claims to support > 2GB processes. We have some ideas on how to get parallelism using multiple CPUs sharing a common memory image (since we can't afford multiple processes each with order 2GB memory).

However, it depends on system V type copy on write fork semantics to allow most of the address space which is effectively r/o to be shared. How similar is OSF in this regard (since the DEC machine is OSF based)? From your knowledge of the OSF release do you know for sure if the >2GB process promise is real for the DEC port?

Thanks
Tim

From: hopper (Mark Hofmann)
Sent: Tuesday, June 27, 1995 3:15 PM
To: tbr (Tim B. Robinson)
Subject: alpha data

Hi Tim,

Unfortunately, I was not able to get all my errands done this afternoon. I may be in late tomorrow morning so I can finish up down here.

I wanted to give you feedback on two items:

1. Synopsys: (from Brian)

Cost/license

duration of license (ie 1 month, year, multi year)

Well, I believe that the latest cronus schedule runs to the beginning of September. So, if we are going to temporarily increase our licenses, we would need them for at least 3-4 months. However, assuming we will have another CMOS project, we will need more than one license for much longer, unless we decide to roll our own.

quick & dirty - based on a run rate of \$3M/month - it's worth a lot to cut the critical path down

I estimate about 33 subblocks. Each subblock takes from an hour to multiple days to map. Assuming the average is about half a day to map one subblock, gives 16-17 license-days for each mapping iteration. Expect a few iterations and resets.

Bottom line: we need to get more. My only question is how many and how quickly we can get them here.

>From one perspective, I count 5-6 users - 3 logic designers, Kleanthes, me, possibly Drew and whatever is in the chipq.

I spoke with Drew yesterday. He counts 3 users (he's not counting himself). I also spoke some more with Brian. By not doing detailed timing models (if we want to just route a functional, low-speed chip for example) we could greatly improve the loop through Synopsys. This might help with the current version of Cronus. However, if we want to re-architect a chip in CMOS we may want more sophisticated models. Clearly there is a trade off between model complexity and runtime. I think 3 Synopsys licenses could be justified if we did a new CMOS architected chip.

2. Alpha box:

hopper writes:

>Dave, Tom, Kurt-

>

> Assume that a single-threaded Dec box is 4x a single-threaded Trex
>and that we can put equal number of processors and memory on both (but
>that Dec has a true 64-bit OS). By how much do you think an Alpha would
>speed up wallclock time in our synthesis, DRC and fracture runs? How
>much time to port stuff over? (Of course, we can point out the binary
>compatibility mode)

>From vanthof Tue Jun 27 06:47:19 1995:

Well, I believe current estimates are between 6-8 weeks, probably closer to 6 weeks, of clock time to run the current synthesis/fracture/drc through on trex. If the alpha is

truely 4x the throughput, then 6 weeks becomes 1.5.

And if tom get's any sort of parallelism done, then we're down to days for a tapeout!

In addition, we don't know for sure that all data will fit within a 2GB memory model. I think it will, but until we get the backend finished, we won't know for sure. Therefore the 64 bit OS is a huge win in it's own right.

I doubt porting vlsimm over to the alpha will be much work, at least not for me :-)

Then there are the multiple number of chips we are going to be working on, each taking a large amount of time:

- euterpe
- mnemosyne
- callioped or/and 2?
- pollux
- castor?
- cronus for csm and tsmc
- mobimos cmos version of cronus

These will all be in the pipeline during the next few months. We need a fast cpu to turn this many chips out.

Dave

>From wampler Tue Jun 27 08:04:53 1995

Ditto to all the points Dave mentioned. Also, as complicated as our synthesis has become, we will undoubtedly need to run some of the layers several times, with full-chip data, to get them fully debugged and producing geometry that is both DRC-clean and in line with the fab folks' style guidelines. Having turnarounds measured in hours or days instead of weeks would give us the ability to iterate when necessary. The current schedule didn't allow for any iterations during fracture, did it (meaning we have to get it right the first time)?

Kurt is correct- It is likely that the actual improvement will be larger as we will need to re-run various parts of the flow.

>From tom Tue Jun 27 18:14:13 1995:

I agree with Dave's calculations. I expect to be able to port vlsimm (the part I'm responsible for, which excludes the mebes read/write and the maskout OPC calculation stuff) in half a week at most. I won't make any predictions about how long it'll take Kurt to port his part, but I wouldn't think it would take long.

>From wampler Tue Jun 27 19:47:52 1995:

I expect the mebesout/mebesread stuff would be an easy port. And gcc is available for the alpha machines (the axposf machine that I have an account on has gcc-2.6.1, alpha-dec-osf1-gcc, and alpha-dec-osf3.0-gcc binaries on it).

I would think I could have it done in a day. The part most likely to require extra work is the (non-essential) UDP broadcast stuff, which resides in a separate library, and need not be linked in immediately if there are problems porting it. The basic mebesout/mebesread stuff is straightforward and should compile with very few changes.

So it looks like the Alpha would be a win.

In addition I would point out that because of the binary compatibilty mode we should be able to run our floating license stuff (Verilog and Gards) faster on the Alpha box!

-hopper

From: paulp (Paul Poenisch)
Sent: Tuesday, June 27, 1995 11:13 AM
To: Tim B. Robinson
Cc: al (Albert Matthews); anh (Anh Ngo); geert (Geert Rosseel); vanthof (Dave Van't Hof)
Subject: Re: Pollux tapeout

tbr wrote:

>
> I have to say I still agree with geert's position on this that with
> 10's of millions of rectangles on one of these chips doing it manually
> is completely ridiculous and the fact that the rules need to be so
> complex that we can't even write them down is going to be a major
> impediment to doing anything quickly and getting it right first time.
> This in turn will be a major competitive disadvantage in the long run
> and will tend to nullify a significant part of the lead we claim to
> have in the technology.
>
> Tim
>

I think that the level of review that I'm talking about is being miss- understood. We do not want to review every rectangle. What we want to do is to review the general layout and floorplan. In particular what we need to see is how the interfaces between particular blocks are handled, how areas that have unusually low utilization and areas that build physical as well as electrical structures. In total we are talking about a dozen or so plots of different, representative areas of the die at 1000 to 5000 X. Also we are concerned about only certain of the layers (about 12) not all of the layers.

As for the feeling that if the rules can not be written down they are too complex, I think that a lot of people are putting far too much faith in computer checks. There is no doubt that computer checks are more reliable than human checks in catching predictable errors, however computers can not catch the problems that no one anticipated. There are two reasons that we can not anticipate all problems with layouts in semiconductor processes. First I am a process engineer and there is no way that I can anticipate every possible layout that a designer might come up with to build something I never thought would show up on one of our IC's. Second design rule checks are based on a model of how the IC process works, and just as the spice model is not completely accurate in all regions, the process model is not accurate for many structures that could be layedout. With the spice model the designers know about not trusting the model in certain operation regions, for the process model the process engineers know not to trust the model used for writing design rules for certain types of structures. The difference between the two is that the process model is orders of magnitude more complex than the spice model and describing all the problem structures in writing is just not practical.

For these two reasons we feel that it is important to visually check layouts for unanticipated problems. As an example, no one in the fab was aware that there were spiral inductors on the first Pollux until we put metal 3 on a wafer, if we had been we could have told the designers early on that the particular design that they had layed out would probably not work (several of our process people have had experience with spiral inductors previously).

Because I never anticipated that anyone would put spiral inductors on one of our IC's I didn't put anything in the rules about them. If I had included them it is likely that it still wouldn't have been right because I have never delt with them before and I may not have talked with the right people to understand the problems before I put the rules out.

I think the argument here is a philosophical one; can the design check process be completely automated. One side says, yes and we should be set up that way from the first day. The other side says, no manual check will always be needed. If we go with the first side the down side risk is months of delay when a design we're counting on fails to make it through the fab. If we go with the second side the down side risk is a few plots and days wasted. I think it's clear which is the more prudent course, if I'm wrong and there are no problems with the layouts then we can forgo the layout reviews in the future, if

I'm right we might all avoid filling out unemployment forms.

Paul.

Subject: Pollux

From: tbr
Sent: Monday, June 26, 1995 10:52 PM
To: paulp (Paul Poenisch)
Cc: al (Albert Matthews); anh (Anh Ngo); Geert Rosseel; vanthof (Dave Van't Hof)
Subject: Re: Pollux tapeout

Paul Poenisch wrote (on Mon Jun 26) :

I don't recall that the Monday meeting said anything about running the DRC flow being SUFFICIENT for the Pollux tape out. It was agreed that we would use the DRC with just a few modifications rather than the trying to apply all of the changes that have been discussed.

I will relay your comments about checking the Pollux layout data on Compass to Al and Anh but since, with the possible exception of Fung, I am the only one from the fab organization who has a clue about using compass I don't think they will be very receptive about doing it that way.

As a general policy I think that a layout review needs to be placed in the normal tapout flow so that this sort of problem does not occur again. Also the review should be done earlier in the process so that it won't be as disruptive to the schedule, certainly before the whole chip DRC and LVS, maby earlier. In any case layout review by the fab is a requirement for any tapeout because, as I've said before, it is impossible to write completely comprehensive design rules, a human visual check is needed to catch those problems that were not anticipated by the rules.

I have to say I still agree with geert's position on this that with 10's of millions of rectangles on one of these chips doing it manually is completely ridiculous and the fact that the rules need to be so complex that we can't even write them down is going to be a major impediment to doing anything quickly and getting it right first time. This in turn will be a major competitive disadvantage in the long run and will tend to nullify a significant part of the lead we claim to have in the technology.

Tim

From: paulp (Paul Poenisch)
Sent: Monday, June 26, 1995 10:38 AM
To: Geert Rosseel
Cc: vanthof (Dave Van't Hof); al (Albert Matthews); anh (Anh Ngo); tbr (Tim B. Robinson)
Subject: Re: Pollux tapeout

>
>
>
> Hi,
>
>
>
> We had a meeting on Monday during wich we decided on a final DRC
> flow. We had agreed that this flow was going to be NECESSARY and
> SUFFICIENT for tape-out of Pollux.
>
> We have decided to tape out Pollux using this DRC flow.
>
>
> A large group "meeting" to review layouts is a waste of time.
> All the layouts are checked in and on-line. There is currently a
> pollux.ly file in the database. If you want to see the layout of
> pollux, it's in
>
> /u/chip/pollux/compass/baseplate/pollux.ly
>
> If you see any structures that you can't make in the fab, we will
> review them and may add them to the Pollux tape-out flow depending on
> how many modules are affected.
>
>
> Geert
>
Hi Geert,

I don't recall that the Monday meeting said anything about running the DRC flow being SUFFICIENT for the Pollux tape out. It was agreed that we would use the DRC with just a few modifications rather than the trying to apply all of the changes that have been discussed.

I will relay your comments about checking the Pollux layout data on Compass to Al and Anh but since, with the possible exception of Fung, I am the only one from the fab organization who has a clue about using compass I don't think they will be very receptive about doing it that way.

As a general policy I think that a layout review needs to be placed in the normal tapout flow so that this sort of problem does not occur again. Also the review should be done earlier in the process so that it won't be as disruptive to the schedule, certainly before the whole chip DRC and LVS, maby earlier.

In any case layout review by the fab is a requirement for any tapeout because, as I've said before, it is impossible to write completely comprehensive design rules, a human visual check is needed to catch those problems that were not anticipated by the rules.

Paul.

Adair L. Alder

From: geert (Geert Rosseel)
Sent: Friday, June 23, 1995 8:24 PM
To: pollux
Subject: Re : layout reviews

Hi Paul,

We had a meeting on Monday during which we decided on a final DRC flow. We had agreed that this flow was going to be NECESSARY and SUFFICIENT for tape-out of Pollux.

We have decided to tape out Pollux using this DRC flow.

A large group "meeting" to review layouts is a waste of time. All the layouts are checked in and on-line. There is currently a pollux.ly file in the database. If you want to see the layout of pollux, it's in

/u/chip/pollux/compass/baseplate/pollux.ly

If you see any structures that you can't make in the fab, we will review them and may add them to the Pollux tape-out flow depending on how many modules are affected.

Geert

DISPATCH AREA

From: vanthof (vant)
Sent: Friday, June 23, 1995 7:26 PM
To: Lisa Robinson
Cc: pollux
Subject: Re: Revised pollux tapeout

Lisa Robinson writes:

>
>
>The following has been revised based upon input from the designers
>
>Lisa Robinson wrote (on Fri Jun 23):

Lisa, it would be interesting to understand why the schedule was pushed out to thursday for meeting the drc checks. I understand the gtlb will need a bunch of work, but I had thought the other modules were quite fixable by wednesday.

Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100

"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

From: vanthof (vant)
Sent: Friday, June 23, 1995 7:24 PM
To: Paul Poenisch
Cc: lisar@microunity.com; pollux
Subject: Re: Pollux tapeout

Paul Poenisch writes:

>
>I hate to mention this at this point but I have yet to see or hear
>anybody mention the layout review for Pollux. If there is something on
>this design which we can't make in the fab all these other check will not matter.
>
>I saw an e-mail about a Pollux review on Monday, is this suppose to be
>the layout review or will be need to schedule it separatly?
>
>Paul.
>

Is this really required? My understanding of the drc situation is that if the chip passes the current set of rules, then it's manufacturable. It may not 'yield' (whatever that means today) very much, but it is manufacturable in the fab.

Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100

"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

From: paulp (Paul Poenisch)
Sent: Friday, June 23, 1995 6:57 PM
To: Lisa Robinson
Cc: pollux
Subject: Re: Pollux tapeout

I hate to mention this at this point but I have yet to see or hear anybody mention the layout review for Pollux. If there is something on this design which we can't make in the fab all these other checks will not matter.

I saw an e-mail about a Pollux review on Monday, is this suppose to be the layout review or will we need to schedule it separately?

Paul.

Arthur L. Owen

take between 6 and 8 weeks of trex CPU time.

Have I missed or mistaken anything?

Lisa R.

From: lisar (Lisa Robinson)
Sent: Friday, June 23, 1995 5:43 PM
To: pollux
Subject: Pollux tapeout

As discussed the following is a summary of the tasks needed to be performed for pollux tapeout. I have underlined the person responsible for completion of each task.

Also below is a list of each module together with the name of the individual responsible for getting the module LVS/DRC clean and where appropriate csyn and celltest clean. All modules should be fully clean by Wednesday. (This may be tight for the gtlb). Vant will work this weekend on the gtlb, getting many of the "minor"

errors fixed, so that eldred can pick up from there on his return on Monday.

mod1 hessam
mod2 dane
mod3 rich
mod4 dane
mod5 hessam
mod6 arya
mod7 rich
mod8 eldred
mod9 geert
mod10 geert
mod11 geert
mod12 geert
scribe mudge(solo)

Tom will have locally run (hopefully to completion - 12 do not route

today) a leafmold rebuild of all of the leaf cells by Monday evening.
This will be picked up Tuesday by tbr, together with as many as

possible of the layout fixes required for pollux, and incorporated in the euterpe/proteus snapshot.

NOTE Timing will NOT be regenerated.

A tar of this snapshot will be taken Wednesday by tom for use as the pollux
--- snapshot. On completion of this,

lisar will start to build a pollux

snapshot to hand over to geert and vant will start up a toplevel

LVS/DRC against this over the long weekend.

Lower layers should be ready ship on Monday July 10, upper layers require back end flow.

Tom and vant estimate end of week beginning

July 10 for the flow. Each module will be synthesized and fractured separately (to maximize parallelism) and the full die will be reassembled at the mask vendor - wampler.

----- In parallel there will be a full chip synthesis
to flush out the process ready for euterpe which cannot be split, this is expected to take
between 6 and 8 weeks of trex CPU time.

Have I missed or mistaken anything?

Lisa R.

From: tom (Tom Laidig [tau])
Sent: Friday, June 23, 1995 12:48 PM
To: hardheads
Cc: tau
Subject: new vlsimm installed

I just installed a new version of vlsimm that contains some new features needed for our tapeout synthesis. The new version has been tested fairly extensively from my home directory, so I don't anticipate any problems. If you think you see one, let me know.

--

TL

Subject: Cronus

From: jack (Jack Wenstrand)
Sent: Thursday, June 22, 1995 7:54 PM
To: cronus
Subject: Foundry notes on the WEB

A much expanded web page has been created for the Cronus foundry work. It is located at
chip -> cronus -> cronus -> process/foundry.

Please check it out, and let me know if keeping something like this up-to-date provides value.

Here are some highlights:

TSMC was here to visit 6/19/95. Minutes are included.

- If all goes well, tape-out to wafers in our hands will take about 33 days.
- 1.6mA/um is a working electromigration limit.
- A 0.5um, 3 level metal, 3.3V process should be available for prototypes by October of this year!

Jack received TSMC 0.5um 3LM 3.3V and 5V design rules and spice models on 6/21. Please let me know if you need a copy.

CSM was here to visit 6/6/95. Minutes follow.

CSM will not run our parts on epi wafers.

We will tape out to both CSM and TSMC in GDS II format. A trial to see how Photronics can handle our files is under way.

Regards,
Jack

Anchor 1.0.0.0

From: lisar (Lisa Robinson)
Sent: Wednesday, June 21, 1995 5:44 PM
To: geert; drew; hopper; tbr
Cc: manser; mudge
Subject: cronus tapeout schedule

Lets meet at 1.00pm tomorrow to discuss what remains to be done for Cronus tapeout.
We'll meet in the boxers conference room.

Lisa R.

Andrew L. Asha

From: tbr
Sent: Tuesday, June 20, 1995 10:24 PM
To: jeff (tfk)
Cc: andrew (Andrew Rhind); billz (Bill Zuravleff); euterpe; jeffm (Jeff Marr); lisar (Lisa Robinson); tfk
Subject: Euterpe Testing (was: DR/DRIO ratio=4)

tfk wrote (on Tue Jun 20):

Our initial plan is to test the function of the synchronizer FIFOs separately, then do functional chip testing at 1:1 or 2:1 clocking.

How?

This should give us decent node coverage but doesn't do any design checking for operation at odd ratios. Certainly we hope that if there is a way to enforce deterministic operation by design, it would be implemented in anything we tape out. Testing non-deterministic hardware is a pain.

It's a fundamental problem in anything which uses multiple clocks (ie is not fully synchronous) since there are synchronization boundaries.
In our case we have 3, the main sofa, the hermes channels, and Cerberus.

I agree that for node coverage you'll get close to 100% coverage at any ratio, but there is the important issue of characterization (ie making sure we meet the goal of being able to absorb phase shift on the Hermes clock). We have tried to do a comprehensive job in simulation, but of course we have no model of metastability in the logic simulator, and while the design is meant to be robust against metastability (it is only sensitive during the initialization phase, even then it will self correct after an error, and in later operation there is no sensitivity) there could be a problem we have not seen in simulation.

Tim

From: jeff (tfk)
Sent: Tuesday, June 20, 1995 8:02 PM
To: lisar (Lisa Robinson); andrew (Andrew Rhind); billz (Bill Zuravleff); euterpe; jeffm (Jeff Marr)
Cc: tfk
Subject: Euterpe Testing (was: DR/DRIO ratio=4)

>Lisa Robinson writes:

>
> Jeff Marr wrote (on Mon Jun 19):
>
> I am only considering ratios that are divisible by 4. A ratio like 5
> would add 20 degrees of uncertainty to the initial state - thus 20
> times more patterns would be required to determine which state
> Euterpe reset in. A ratio of 4 has zero degrees of uncertainty; a
> ratio of 8 has one. Thus a ratio of 8 would double the number of
> patterns required. All these doublings, and 20 time-sings are
> relative to the ratio of 4.
>
> Doubling what, you say? The set of patterns for determining the
> hermes channel state.
>
> I am not sure how many patterns will be required - hopefully, we can
> find a way to keep it to 32 to 48 per channel (4 orders of
> uncertainty from the relationship of pll0 and pll1 outputs, and 8 to
> 12 orders of uncertainty because of the hermes channel enable's
> relationship to internal counters). For two channels, because of
> shared pll1 output, that would be 256 to 576 patterns. Having to
> settle for a dram ratio of 8 would double this range.
>
> jeffm
>
> Jeff, Roughly, how many vectors (or minor cycles) are being doubled?
> Is the synchronization done once at the beginning of the test or for
> each page.
>
> Lisa R.
>
>Synchronization will be done at the beginning of the test only,
>provided that the tester can assure that the number of breakvectors is
>a multiple of a number that we can specify (such as 64) (comment,
>Andrew?).

Number of break vectors can be determined absolutely after a MATCH which is what I think you are referring to here. A MATCH instruction causes the tester to loop through a set of vectors until the DUT outputs match the tester's expected data.

Limitation:

The looping cycle count is modulo 16 (or 32, I'll check), so if the match pattern from the DUT is also an integer power of two (with some exceptions) then the patterns might never match up.

>
>Based on the talk that I just had with tbr, here are some rough
>numbers:
>
> If the hermes channels are run at a rate equal to the sofa
> rate, then the degrees of uncertainty are pretty highly
> constrained. There are basically 4 degrees of uncertainty -

> since it is initially not known how the hermes return data will
> line up with eta.

I'm not sure I understand this. Do you mean that the data could show up on any one of 4 output clocks relative to input hermes data or that the data itself might have 4 different sets of contents? (the first is relatively easy to deal with, the second is nearly impossible).

> The uncertainty in the outgoing hermes data is in the phase
> relationship between it in the reference clock - supposedly the
> tester can compensate (comment, Andrew?).

The tester can certainly have it's timing altered to match the DUT. The only tool that the tester provides to have this happen automatically works only with an external clock driving the tester.

In that case it can move it's timing around to find a designated transistion. (This is the tester's SYNC command) To do this for internal clocking or to use more sophistication requires externally created utilities manipulating the tester. These do not currently exist.

> If the hermes channels are run at a rate different than the
> sofa, then the degrees of uncertainty go up a lot. For
> example, if the multiplier ratios are 11/10, then there are $2 * 11 * 4$ degrees of uncertainty (The frame and eta, as well as
> the uncertainty in the rate fifo, interact to give 88
> possibilities). With precise phase control on the returning idle
> packets this may be reduced to 44, since the two input fifos
> could be coerced to have the same separation in the input and
> output pointers.

> Assuming that the hermes channels are run at the sofa rate,
> then, as a very rough estimate, the synch pattern would consist
> of 2 hermes operations and two dram operations - say about 450
> vectors. THIS IS A ROUGH GUESS! It could take more operations
> to be able to tell the potentially subtle difference between
> the 4 hermes input cases. I believe that if one channel is
> synched the other is synched, but only in the case where the
> hermes rate equals the sofa rate. If the dram ratio is 4, then
> that means about 1800 vectors are required, but if it is 8 then
> that translates to 3600 vectors.

>The tightly constrained vanilla case is quite reasonable, but there is a
>very large (2 to 3, or more, orders of magnitude) step function if
>different interface timing is desired. Is different timing desired?

It depends on how thoroughly we wish to test the chip. Matching to different vector sets is very involved on this tester. There is no built-in conditional branching (besides MATCH and SYNC, which are only go/no-go) so this would all have to be controlled external from the tester by downloading error maps, comparing results then reloading new vector sets. Likely incompatible with any attempt to test quickly and also quite inflexible as tests change.

Our initial plan is to test the function of the synchronizer FIFOs separately, then do functional chip testing at 1:1 or 2:1 clocking. This should give us decent node coverage but doesn't do any design checking for operation at odd ratios. Certainly we hope that if there is a way to enforce deterministic operation by design, it would be implemented in anything we tape out. Testing non-deterministic hardware is a pain.

--

tfk

tfk@microunity.com

Sunnyvale, CA

How come the dove gets to be the peace symbol? How about the pillow?
It has more feathers than the dove, and it doesn't have that dangerous beak.

From: tbr
Sent: Tuesday, June 20, 1995 1:18 PM
To: hopper (Mark Hofmann)
Cc: Geert Rosseel; vant
Subject: Re: DRC flow

Mark Hofmann wrote (on Tue Jun 20):

Geert Rosseel writes:
Hi Dave,

When you have the DRC flow that incorporates all the latest changes that wediscussed yesterday, can you send out a mail to hard-heads to let everyone know that we have frozen the DRC flow for Euterpe & Pollux tape-outs ?

Thank's

Geert,

Have you and Tim been discussing this situation with Steve Manser?
I'll be back in tomorrow.

Yes, and mouss is in the loop.

Tim

Amber L. Allen

From: vanthof (vant)
Sent: Tuesday, June 20, 1995 12:43 PM
To: Bruce Bateman
Cc: geert (Geert Rosseel); tbr (Tim B. Robinson); hopper (Mark Hofmann); tom (Tom Laidig); vanthof (Dave Van't Hof)
Subject: Re: frozen drc flow

Bruce Bateman writes:

>
>Have we made a decision yet as to whether to put the cache memory cell
>changes into the euterpe tape out? The "altered" cell is still sitting
>in my local directory and I have not done any further work on it since
>the original review and "drc" run to the old flow.
>
>BB
>

Well, I believe the changes are what the fab wants, so I'd say yes, we should use it. The current drc flow does not yet address the false errors we get from the cache cell. I'm going to work on those next. The main intent of the current flow was to start checking the new rules.

Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

Amber L. Adkins

From: stick (Bruce Bateman)
Sent: Tuesday, June 20, 1995 12:13 PM
To: geert; tbr
Cc: vanthof; hopper; tom
Subject: Re: frozen drc flow

> From: vanthof (vant)
> Subject: frozen drc flow
> To: hardheads
> Date: Tue, 20 Jun 95 10:00:15 PDT
> Cc: vanthof (Dave Van't Hof)
>
> With the latest release of the drc flow (last night), the drc rules
> have been frozen for the euterpe and pollux tapeouts. Please verify
> any layouts used in these chips against this drc flow.
>
> Thanks,
> Dave

Have we made a decision yet as to whether to put the cache memory cell changes into the euterpe tape out? The "altered" cell is still sitting in my local directory and I have not done any further work on it since the original review and "drc" run to the old flow.

BB

Subject: [REDACTED]

From: vanthof (vant)
Sent: Tuesday, June 20, 1995 12:00 PM
To: hardheads
Cc: vanthof (Dave Van't Hof)
Subject: frozen drc flow

With the latest release of the drc flow (last night), the drc rules have been frozen for the euterpe and pollux tapeouts. Please verify any layouts used in these chips against this drc flow.

Thanks,
Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100

"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

From: vanthof (vant)
Sent: Tuesday, June 20, 1995 11:57 AM
To: Geert Rosseel
Cc: geert@microunity.com; hopper (Mark Hofmann); tbr (Tim B. Robinson); vant
Subject: Re: DRC flow

Geert Rosseel writes:

>
>> don asbestos suits...
>
> i think we got to some agreement yesterday with paul on what we're
>going to fix and what we're going to leave for next time ...
>
> There is still some major contention on how we're going to deal with
>the pad-rules.
>
> Geert

Yes, that was what I gathered from the meeting yesterday as well.

I'll just send out a simple note stating that the drc flow is now frozen for the euterpe and pollux tapeouts.

Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just
not in this context." The Tick to Thrackazog

From: lisar (Lisa Robinson)
Sent: Tuesday, June 20, 1995 11:55 AM
To: tbr
Subject: forwarded message from Tom Sanders

----- Start of forwarded message -----

Return-Path: <toms>
Received: from hera.microunity.com by gaea.microunity.com (4.1/musel.3)
id AA12198; Tue, 20 Jun 95 09:45:16 PDT
Received: by hera.microunity.com (8.6.10/muse-sw.3)
id JAA08595; Tue, 20 Jun 1995 09:45:15 -0700
Message-Id: <199506201645.JAA08595@hera.microunity.com>
X-Status: N
X-Mailer: Applixware 3.1(473)
From: toms (Tom Sanders)
To: jeffm, lisar, jeff, andrew
Cc: mudge
Subject: Euterpe synchronization
Date: Tue, 20 Jun 1995 09:45:15 -0700

Can we meet at 2:00 on Wednesday in pecr to discuss the hermes channel synch issue on euterpe. We need a better understanding of the issues involved in order to determine a solution using the tester.

----- End of forwarded message -----

Amber L. Akins

From: hopper (Mark Hofmann)
Sent: Tuesday, June 20, 1995 4:18 AM
To: Geert Rosseel
Cc: tbr (Tim B. Robinson); vant
Subject: Re: DRC flow

Geert Rosseel writes:

Hi Dave,

When you have the DRC flow that incorporates all the latest changes that wediscussed yesterday, can you send out a mail to hard-heads to let everyone know that we have frozen the DRC flow for Euterpe & Pollux tape-outs ?

Geert-

I think it's not a bad idea to do this. I think it's a good idea. But since this might be inflammatory, maybe we want to make sure how this note is written and be prepared to don asbestos suits...

-hopper

From: hopper (Mark Hofmann)
Sent: Tuesday, June 20, 1995 3:59 AM
To: Geert Rosseel
Cc: vant; tbr (Tim B. Robinson)
Subject: Re: DRC flow

Geert Rosseel writes:

Hi Dave,

When you have the DRC flow that incorporates all the latest changes that wediscussed yesterday, can you send out a mail to hard-heads to let everyone know that we have frozen the DRC flow for Euterpe & Pollux tape-outs ?

Thank's

Geert,

Have you and Tim been discussing this situation with Steve Manser?
I'll be back in tomorrow.

-hopper

From: geert (Geert Rosseel)
Sent: Tuesday, June 20, 1995 10:56 AM
To: vanthof
Cc: hopper; paulp; tau; tbr
Subject: DRC flow

Hi Dave,

When you have the DRC flow that incorporates all the latest changes that wediscussed yesterday, can you send out a mail to hard-heads to let everyone know that we have frozen the DRC flow for Euterpe & Pollux tape-outs ?

Thank's

Printed on: 6/20/95 10:11 AM

From: vanthof (vant)
Sent: Tuesday, June 20, 1995 9:11 AM
To: Mark Hofmann
Cc: vnathof; geert (Geert Rosseel); tbr (Tim B. Robinson); tom (Tom Laidig)
Subject: Re: Ringoscillator tape-out

Mark Hofmann writes:

>
>Okay. I keep forgetting that. Should we be able to avoid holes if we
>only add routing over sofa atoms?
>

I believe so, but Tom is the router expert, so he would know for sure.
Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100

"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

From: hopper (Mark Hofmann)
Sent: Tuesday, June 20, 1995 2:07 AM
To: vant
Cc: hopper@microunity.com; geert (Geert Rosseel); vant; tbr (Tim B. Robinson); tau
Subject: Re: Ringoscillator tape-out

vant writes:

Mark Hofmann writes:

>Do these sdec-iso problems stop (functionally) devices in the ECL sofa
>atom (non-end cells) from working? If these problems are not show stoppers
>then we could replicate lots of small odd-numbered (say 3)
inverters chains
>on the ECL sofa. Some should work.

I'm trying to remember just what Al drew on the board, but I think it had
to do with shorts in the sdec layer. The current sdec-iso synthesis fixed
the problems it was addressing, but there were cases it missed. The new
sdec-iso synthesis is totally in the back end should not require
any additional
work to tapeout, just a bit of cpu time. I think it would be a good idea
to tapeout sdec-iso out as well.

Okay. Make sense. It's just an extra mask which gives us some insurance then.

>Could we route these by hand in wide metal? I'm thinking of making
>everything as "non-minimum" as possible for best yield.

We could also tell the router to use non-minimum wires if we wanted.
At least I think we could.

> If you don't care about the lower layers and we are careful about
what atoms
> we use, then we punt on drc's like you said. I think we want some form
> of lvs to make sure that we don't screw something simple up and tapeout
> a rock.
>
>Agreed. A simple LVS flow is a good idea (Would our current LVS work?)

It should. The lvs flow hasn't changed much. It may take some funky
windowing to hide some of the layouts with unused devices (and to make it
run faster), but I've done that stuff before.

Great. Sounds like that will work.

>Can we create this (with fat metals and so on) so we're sure we don't have
>any notches?

It's not really notches that are the problem, it's holes. The notch rule
was a simplification of the hole rule to make it checkable. So even if we
create notches, it should be okay in the fab. We should just avoid making
holes less than 1.5 microns.

Okay. I keep forgetting that. Should we be able to avoid holes if we only add routing over
sofa atoms?

Dave

-hopper

From: vanthof (vant)
Sent: Tuesday, June 20, 1995 8:55 AM
To: Mark Hofmann
Cc: geert (Geert Rosseel); vant; tbr (Tim B. Robinson); tau
Subject: Re: Ringoscillator tape-out

Mark Hofmann writes:

>
>Do these sdec-iso problems stop (functionally) devices in the ECL sofa
>atom (non-end cells) from working? If these problems are not show
>stoppers then we could replicate lots of small odd-numbered (say 3)
>inverters chains on the ECL sofa. Some should work.

I'm trying to remember just what Al drew on the board, but I think it had to do with shorts in the sdec layer. The current sdec-iso synthesis fixed the problems it was addressing, but there were cases it missed. The new sdec-iso synthesis is totally in the back end should not require any additional work to tapeout, just a bit of cpu time. I think it would be a good idea to tapeout sdec-iso out as well.

>
>Could we route these by hand in wide metal? I'm thinking of making
>everything as "non-minimum" as possible for best yield.

We could also tell the router to use non-minimum wires if we wanted. At least I think we could.

>
> If you don't care about the lower layers and we are careful about what atoms
> we use, then we punt on drc's like you said. I think we want some form
> of lvs to make sure that we don't screw something simple up and tapeout
> a rock.
>
>Agreed. A simple LVS flow is a good idea (Would our current LVS work?)

It should. The lvs flow hasn't changed much. It may take some funky windowing to hide some of the layouts with unused devices (and to make it run faster), but I've done that stuff before.

>
>Can we create this (with fat metals and so on) so we're sure we don't
>have any notches?

It's not really notches that are the problem, it's holes. The notch rule was a simplification of the hole rule to make it checkable. So even if we create notches, it should be okay in the fab. We should just avoid making holes less than 1.5 microns.

Dave

--

Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100

"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

From: hopper (Mark Hofmann)
Sent: Tuesday, June 20, 1995 1:35 AM
To: geert (Geert Rosseel); vant; tbr (Tim B. Robinson); tau
Subject: Re: Ringoscillator tape-out

vant writes:

Geert Rosseel writes:

> There seems to be a general agreement that we need a moral boost
> in the company to keep people going at this pace ..

Yes, that would be really good.

> I was thinking about what we could do to get something
> out of the fab. As I understand from Al and Anh, the
> process is working quite well up to M2 . M3 and M4 still
> need work.

>
> So, I was thinking about what is the simplest thing
> we could do to get a working something ...

>
> Apparently, we have a good calliope I baseplate in the fab
> ready for metal processing.

We'd have to be careful about what atoms in the sofa we were to use.
The min diffusion problem pointed out last week affects (I believe) only
atoms at each end of each row. by not using those atoms, we should be
able to get working logic blocks. Then there is the question about biasing,
clocking, etc. Then there has been a request for a new backend synthesis
for sdec-iso to fix some other problems Al has found.

Do these sdec-iso problems stop (functionally) devices in the ECL sofa atom (non-end
cells) from working? If these problems are not show stoppers then we could replicate lots
of _small_ odd-numbered (say 3) inverters chains on the ECL sofa. Some should work.

> The idea is to manually route one or more oscillators
> in M1 on SOFA (CMOS or ECL).

>
> I think we can get by with 4 masks. We can build and fab these
> in about a week. The actual lay-out is probably also a week ...

It may be more than 4 masks if we really want something to work
and be testable. The additional masks should not add that much to the
overall schedule.

> So, we might have somethin working in 2 weeks ... I would
> like to go for this ...

>
> I think we may have to punt on DRC and LVS for this ...

with moderate effort, I think we could have something relatively drc clean
IN THE SOFA REGION, but that would mean a full mask set upto metal 1 or
metal 2 (wherever you want to stop).

Could we route these _by hand_ in wide metal? I'm thinking of making everything as "non-
minimum" as possible for best yield.

If you don't care about the lower layers and we are careful about what atoms
we use, then we punt on drc's like you said. I think we want some form
of lvs to make sure that we don't screw something simple up and tapeout
a rock.

Agreed. A simple LVS flow is a good idea (Would our current LVS work?)

> Questions :

>

> Can we get a visual picture of Calliope up to SDEC iso, so
> we can draw the extra layers on top of that.

>

> Is it easy to do the back-end flow on this ..

I have a tapeout flow for lower layers ready. The lower metals are almost there and with minimal effort could be ready to do all the backend work except notch filling, which sounds like Paul may want to drop anyway (besides, it's not a functionality problem).

Can we create this (with fat metals and so on) so we're sure we don't have any notches?

> Is there any way we can run LVS/DRC ...

I think so, with some work.

> Can anyone see any major obstacles ????

My concern is the use of the existing baseplate. Sure we can use the sofa region for connecting up some gates, but the ecl sofa requires bias generators and some form of clock. Do we think the min diffusion bug will not affect those features? If we think they will, can we route the required lines to various pads and hookup external sources? If that's the case, then let's do it.

As Dave and Tim mention- the bias generator/reference circuitry adds some complication to the ECL inverter test. Is there a way around this?

Could we also build CMOS ring oscillators? (using the cerberus area?)

Cool stuff.

Dave

-hopper

From: tbr
Sent: Tuesday, June 20, 1995 1:08 AM
To: geert (Geert Rosseel)
Cc: hopper; tau; vanthof
Subject: Ringoscillator tape-out

Geert Rosseel wrote (on Mon Jun 19):

The idea is to manually route one or more oscillators
in M1 on SOFA (CMOS or ECL).

Can this be done without needing th bias infrastructure to be working?

Tim

From: vanthof (vant)
Sent: Tuesday, June 20, 1995 12:39 AM
To: Geert Rosseel
Cc: hopper (Mark Hofmann); tau; tbr (Tim B. Robinson); vanthof (Dave Van't Hof)
Subject: Re: Ringoscillator tape-out

Geert Rosseel writes:

>
>
> Hi,
>
> There seems to be a general agreement that we need a moral boost in
>the company to keep people going at this pace ..

Yes, that would be really good.

>
> I was thinking about what we could do to get something out of the fab.
>As I understand from Al and Anh, the process is working quite well up
>to M2 . M3 and M4 still need work.
>
> So, I was thinking about what is the simplest thing we could do to get
>a working something ...
>
> Apparently, we have a good calliope I baseplate in the fab ready for
>metal processing.

We'd have to be careful about what atoms in the sofa we were to use.
The min diffusion problem pointed out last week affects (I believe) only atoms at each end
of each row. by not using those atoms, we should be able to get working logic blocks.
Then there is the question about biasing, clocking, etc. Then there has been a request
for a new backend synthesis for sdec-iso to fix some other problems Al has found.

>
> The idea is to manually route one or more oscillators in M1 on SOFA
>(CMOS or ECL).
>
> I think we can get by with 4 masks. We can build and fab these in
>about a week. The actual lay-out is probably also a week ...

It may be more than 4 masks if we really want something to work and be testable. The
additional masks should not add that much to the overall schedule.

>
> So, we might have somethin working in 2 weeks ... I would like to go
>for this ...
>
> I think we may have to punt on DRC and LVS for this ...

with moderate effort, I think we could have something relatively drc clean _IN THE SOFA
REGION_, but that would mean a full mask set upto metal 1 or metal 2 (wherever you want to
stop).

If you don't care about the lower layers and we are careful about what atoms we use, then
we punt on drc's like you said. I think we want some form of lvs to make sure that we
don't screw something simple up and tapeout a rock.

>
> Questions :
>
> Can we get a visual picture of Calliope up to SDEC iso, so we can
>draw the extra layers on top of that.

>
> Is it esay to do the back-end flow on this ..

I have a tapeout flow for lower layers ready. The lower metals are almost there and with minimal effort could be ready to do all the backend work except notch filling, which sounds like Paul may want to drop anyway (besides, it's not a functionality problem).

>
> Is there any way we can run LVS/DRC ...

I think so, with some work.

>
> Can anyone see any major obstacles ????

>
My concern is the use of the existing baseplate. Sure we can use the sofa region for connecting up some gates, but the ecl sofa requires bias generators and some form of clock. Do we think the min diffusion bug will not affect those features? If we think they will, can we route the required lines to various pads and hookup external sources? If that's the case, then let's do it.

Cool stuff.
Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100

"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

vanthof (vant)
Monday, June 19, 1995 11:15 AM
andrew
euterpe; vanthof (Dave Van't Hof)
RE: Tapeout status

andrew writes:

>
>Tim
>
>We are working on a Membrane Probe Card for Euterpe and need to supply
>the vendor with pad function and x-y info. Has the Euterpe pad
>definition been finalized? Including all Vdd and Vss pads? Where can
>I get the final list of pad definitions and an x-y location for each
>pad - including all internal pads.
>
>Thanks
>Andrew

I generated a list of this kind for Mark W. a while back. However, the pads have changed and this list must be regenerated. It is not a straight forward task and requires some hand intervention. However, I believe I can modify what I've got with a minimum amount of effort. I don't know what Mark did with what I gave him, but I know it was not in any format that any vendor could use. If you let me know what format file you need, I think I can generate that for you.

Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"I don't know the meaning of the word surrender! I mean, I know it, I'm not dumb... just not in this context." The Tick to Thrackazog

From: tbr
Sent: Sunday, June 18, 1995 9:59 PM
To: tom (Tom Laidig [tau])
Cc: Geert Rosseel; lisar (Lisa Robinson); tau
Subject: Re: cebuf8

Tom Laidig [tau] wrote (on Sun Jun 18):

Geert Rosseel writes:

| Thank's Lisa .. Now I know what it is ...

OK, does this mean I don't need to pull stuff off the pollux snapshot tape? (the online stuff is indeed only the layout data, which was restored piecemeal from tape)

Correct.

Tim

From: tom (Tom Laidig [tau])
Sent: Sunday, June 18, 1995 9:56 PM
To: Geert Rosseel
Cc: lisar (Lisa Robinson); tbr (Tim B. Robinson); tau
Subject: Re: cebuf8

Geert Rosseel writes:

| Thank's Lisa .. Now I know what it is ...

OK, does this mean I don't need to pull stuff off the pollux snapshot tape? (the online stuff is indeed only the layout data, which was restored piecemeal from tape)

--

, _

From: hopper (Mark Hofmann)
Sent: Sunday, June 18, 1995 12:47 PM
To: Geert Rosseel
Cc: al (Albert Matthews); anh (Anh Ngo); paulp (Paul Poenisch); tau; tbr (Tim B. Robinson); vanthof (Dave Van't Hof); manser; mouss (John Moussouris)
Subject: Re: DRC FLOW STRATEGY

Geert Rosseel writes:

I think we have a serious problem on our hands in the physical verification of our designs.

The design group wants a set of machine-checkable design rules such that if a design meets these rules, it is manufacturable.

We have currently a DRC flow that if we meet it, to the best of our knowledge, the design will NOT work. I know of a lot of structures that meet the rules and will not work.

The fab-organization has taken the position that they will visually check layouts for these cases and correct them one by one, which is a rather preposterous proposition considering the number of polygons on our designs.

A good example of this impasse are the pad-structures :

We had pads before that met the DRC rules. They failed rather badly. We are now fixing these pads. Our most senior mask-designer has now spend 2 weeks full-time on these pads and needs at least another week to finish this. These pads are layed out to specification from the fab-organization that are not written down and not checked. More-over there is no guarantee that these structures do not appear anywhere else on the design or another design and we have no way to find out.

This is completely the wrong way to do this.
The right way to do this is as follows :

If you find a structure that does not work

1. You write it down in the DRC document
so that the mask-designers know to avoid these structures in advance.
2. You code the forbidden structure in the DRC flow
so we can find ALL forbidden structures
3. You run the DRC flow on the WHOLE design and fix ALL similar violations.

There is an additional confusion going on on nomenclature and I want to set this straight once and for all :

FUNCTIONALITY : If a design rules affects functionality, it means that this structure just will not work. It means that ALL similar structures have to be fixed on the die to have a functional die.
These are the rules that are defined in the DRC flow.

YIELD : if a design rules affects yield, it means that we can make a functional die, even with a lot of these structures,
however, we can bring the yield up (and make cheaper parts) by following these rules.

this "rule" is not a "rule" but a guideline.

To summarize :

a drc rule	affects functionality
a drc guideline	affects yield

The fab-organization has to decide in which category a certain rule falls.

THEREFORE :

Our current problem is not yield, but functionality. So, we should concentrate on the DRC flow. From now on, we will concentrate on the DRC flow under the following strategy.

Any of our designs will meet the DRC flow and if it does, it is ready for tape-out.

If the above is not true, we will (in this order)

1. fix the flow
2. fix the design to meet the new flow

To get this process started, we will meet 3 times a week (Monday, Wednesday and Friday) at 11:30 for no more than 30 minutes.

Thank's
Geert

Well said.
Let's do it.

-hopper

From: tbr
Sent: Sunday, June 18, 1995 6:12 PM
To: solo
Subject: forwarded message from Geert Rosseel

FYI. Mail from geert . . .

----- Start of forwarded message -----

Status: RO

X-VM-v5-Data: ([nil nil nil nil t nil nil nil nil]

["3233" "Sat" "17" "June" "1995" "11:13:52" "-0700" "Geert Rosseel" "geert "
"<199506171813.LAA23729@ambiorix.microunity.com>" "89" "DRC FLOW STRATEGY" "^From:" nil
nil "6"])

Return-Path: <geert>

Received: from ambiorix.microunity.com by gaea.microunity.com (4.1/muse1.3)

id AA10607; Sat, 17 Jun 95 11:13:53 PDT

Received: by ambiorix.microunity.com (8.6.10/muse-sw.3)

id LAA23729; Sat, 17 Jun 1995 11:13:52 -0700

Message-Id: <199506171813.LAA23729@ambiorix.microunity.com>

From: geert (Geert Rosseel)

To: al, anh, hopper, paulp, tau, tbr, vanthof

Cc: manser, mouss

Subject: DRC FLOW STRATEGY

Date: Sat, 17 Jun 1995 11:13:52 -0700

Hi,

I think we have a serious problem on our hands in the physical verification of our designs.

The design group wants a set of machine-checkable design rules such that if a design meets these rules, it is manufacturable.

We have currently a DRC flow that if we meet it, to the best of our knowledge, the design will NOT work. I know of a lot of structures that meet the rules and will not work.

The fab-organization has taken the position that they will visually check layouts for these cases and correct them one by one, which is a rather preposterous proposition considering the number of polygons on our designs.

A good of example of this impasse are the pad-structures :

We had pads before that met the DRC rules. They failed rather badly. We are now fixing these pads. Our most senior mask-designer has now spend 2 weeks full-time on these pads and needs at least another week to finish this. These pads are layed out to specification from the fab-organization that are not written down and not checked. More-over there is no gaurantee that these structures do not appear anywhere else on the design or another design and we have no way to find out.

This is completely the wrong way to do this.
The right way to do this is as follows :

If you find a structure that does not work

1. You write it down in the DRC document
so that the mask-designers know to avoid these structures
in advance.

2. You code the forbidden structure in the DRC flow
so we can find ALL forbidden structures
3. You run the DRC flow on the WHOLE design and fix ALL similar violations.

There is an additional confusion going on on nomenclature and I want to set this straight once and for all :

FUNCTIONALITY : If a design rules affects functionality, it means that this structure just will not work. It means that ALL similar structures have to be fixed on the die to have a functional die.
These are the rules that are defined in the DRC flow.

YIELD : if a design rules affects yield, it means that we can make a functional die, even with a lot of these structures, however, we can bring the yield up (and make cheaper parts) by following these rules.
this "rule" is not a "rule" but a guideline.

To summarize :

a drc rule	affects functionality
a drc guideline	affects yield

The fab-organization has to decide in which category a certain rule falls.

THEREFORE :

Our current problem is not yield, but functionality. So, we should concentrate on the DRC flow. From now on, we will concentrate on the DRC flow under the following strategy.

Any of our designs will meet the DRC flow and if it does, it is ready for tape-out.
If the above is not true, we will (in this order)

1. fix the flow
2. fix the design to meet the new flow

To get this process started, we will meet 3 times a week (Monday, Wednesday and Friday) at 11:30 for no more than 30 minutes.

Thank's

Geert

----- End of forwarded message -----

Number 1. A-4mp

From: solo (John Campbell)
Sent: Sunday, June 18, 1995 5:12 PM
To: Tim B. Robinson
Subject: Re: Tapeout status

as Tim B. Robinson was saying

..The main gating factor has been cleaning up DRC errors from the latest ..design rule changes, particularly in the area of the pads. These ..changes have now been made but we have not yet confirmed the full chip ..is still clean at the top level (this is about a 4 day run). However, ..the fab has now requested yet more changes and we will be assessing ..the impact of these Monday. We are ready to fracture the lower layers ..just as soon as the rules remain stable long enough to do it.

invite me to your decision meet on the update of snapshot. we have a few cells with remaining errors. more on monday.

....
regards,
solo a.k.a. John Campbell x516

From: tbr
Sent: Sunday, June 18, 1995 1:59 PM
To: geert (Geert Rosseel)
Cc: billz; dickson; mws; wampler; woody
Subject: Re: Latest route

Geert Rosseel wrote (on Sun Jun 18) :

Hi Tim,

Is all this because of your new routing strategy ? Maybe we really should think of a way to change these +- 100 gates "in place" after the route. If we find a way to do this (even if it's ugly or cumbersome) , we probably should just do it and tape this thing out.

I think I can account for the apparantly better result. The new routing order is basically just ordered by slack time, adjusted for net length. As a result, a lot of longer nets are getting left till later. So I assume the total amount of route in the design is probably the same as before because a higher proportion of the nets are shorter. Looking at the unroutes, the purple is much more uniformly distributed because many of the unroutes are longer nets.

That said however, the timing looks much better for the stuff that is routed.

Tim

From: tbr
Sent: Sunday, June 18, 1995 12:38 PM
To: geert (Geert Rosseel)
Cc: billz; dickson; mws; wampler; woody
Subject: Re: Latest route

Geert Rosseel wrote (on Sun Jun 18):

Hi Tim,

Is all this because of your new routing strategy ? Maybe we really should think of a way to change these +- 100 gates "in place" after the route. If we find a way to do this (even if it's ugly or cumbersome) , we probably should just do it and tape this thing out.

As far as I know, that's the only thing of substance that changed. I did pick up a minor placement change for a bug fix in uu.

Tim

From: geert (Geert Rosseel)
Sent: Sunday, June 18, 1995 12:33 PM
To: tbr
Cc: billz; dickson; mws; wampler; woody
Subject: Re: Latest route

Hi Tim,

Is all this because of your new routing strategy ? Maybe we really should think of a way to change these +- 100 gates "in place" after the route. If we find a way to do this (even if it's ugly or cumbersome) , we probably should just do it and tape this thing out.

Geert

From: geert (Geert Rosseel)
Sent: Saturday, June 17, 1995 1:14 PM
To: al; anh; hopper; paulp; tau; tbr; vanthof
Cc: manser; mouss
Subject: DRC FLOW STRATEGY

Hi,

I think we have a serious problem on our hands in the physical verification of our designs.

The design group wants a set of machine-checkable design rules such that if a design meets these rules, it is manufacturable.

We have currently a DRC flow that if we meet it, to the best of our knowledge, the design will NOT work. I know of a lot of structures that meet the rules and will not work.

The fab-organization has taken the position that they will visually check layouts for these cases and correct them one by one, which is a rather preposterous proposition considering the number of polygons on our designs.

A good example of this impasse are the pad-structures :

We had pads before that met the DRC rules. They failed rather badly. We are now fixing these pads. Our most senior mask-designer has now spend 2 weeks full-time on these pads and needs at least another week to finish this. These pads are laid out to specification from the fab-organization that are not written down and not checked. More-over there is no guarantee that these structures do not appear anywhere else on the design or another design and we have no way to find out.

This is completely the wrong way to do this.
The right way to do this is as follows :

If you find a structure that does not work

1. You write it down in the DRC document
so that the mask-designers know to avoid these structures
in advance.
2. You code the forbidden structure in the DRC flow
so we can find ALL forbidden structures
3. You run the DRC flow on the WHOLE design and fix ALL similar violations.

There is an additional confusion going on on nomenclature and I want to set this straight once and for all :

FUNCTIONALITY : If a design rules affects functionality, it means that this structure just will not work. It means that ALL similar structures have to be fixed on the die to have a functional die.
These are the rules that are defined in the DRC flow.

YIELD : if a design rules affects yield, it means that we can make a functional die, even with a lot of these structures, however, we can bring the yield up (and make cheaper parts) by following these rules.

this "rule" is not a "rule" but a guideline.

To summarize :

a drc rule	affects	functionality
a drc guideline	affects	yield

The fab-organization has to decide in which category a certain rule falls.

THEREFORE :

Our current problem is not yield, but functionality. So, we should concentrate on the DRC flow. From now on, we will concentrate on the DRC flow under the following strategy.

Any of our designs will meet the DRC flow and if it does, it is ready for tape-out. If the above is not true, we will (in this order)

1. fix the flow
2. fix the design to meet the new flow

To get this process started, we will meet 3 times a week (Monday, Wednesday and Friday) at 11:30 for no more than 30 minutes.

Thank's

Geert

From: doi (Derek Iverson)
Sent: Thursday, June 15, 1995 2:54 PM
To: guarino; gmo; doi; gregg; lisar; jeffm; lisar
Cc: hestia
Subject: Software Bringup Meeting Minutes - June 14, 1995

Software Bringup Meeting

June 14, 1995

Next Meeting: June 21 at 2:00 pm.

Note: Loretta and I will both be absent from the meeting
if one is held.

Attendees: guarino, gmo, doi, gregg, lisa, jeffm, lisar

New Action Items

None.

Review of Action Items

Item: Build osf kernel and perpare it for execution on HW simulator

Who: doi/iimura

Status: Pending

06/14 Derek is going to send e-mail to lisar on how to accomplish this.

Item: Plan for testing remote debugging environment

Who: everyone

Status: Pending

04/13 This will be discussed at the next meeting.

04/20 Short discussion on the pieces required and if we want to have
a standalone version of the hostio software.

Wally has started on this already.

04/26 Any work with snoopy/dram models will begin after euterpe tapeout
when jeff has more time.

05/03 Wayne F. expects to have a prototype ISA card ready for testing
tomorrow. Gmo is in the process of writing code to test the card.

05/10 Gmo has written a test program. Gmo and Wayne need to get together
and try it on the board.

05/24 Wally has added more connectivity between the ukernel and the
debugger (ability to interrupt kernel via the cerberus
'forced interrupt' bit).

Gmo is ready with a test when Wayne has the board ready.

05/31 Wayne has card ready to test (fpga that talks to ISA) for initial
testing. Documentation for the interface next week some time.

06/14 Gmo has been able to access Wayne's board over the ISA bus.
Informal review of the ISA CBI device later today.

Item: ukernel needs to detect machine checks and 'do the right thing'

Who: lisa

Status: Pending

04/13 No new progress.
04/20 On list of things to do. Lower priority.

Item: Create performance test plan

Who: claseman

Status: [11/30] No progress as focus is on functionality.

11/30 We continue to run tests to help us compare terp vs hardware performance.
We still need to put together the actual performance tests that need to be run on the hardware.
05/10 Back to life. Tim has checked in more tests. We are going to use the tests written for hermes and cerberus read accesses in the 'How to debug Euterpe' presentation on monday. We will be looking for the time it takes from the instruction issue, entry to nb, and completion of access.
05/24 Tim is going to send lisar a list of the test names that he intends to write so they can be incorporated into the test template.
05/31 Jeffm is going to analyze approx one test per day (with Tim looking over his shoulder) to determine/verify the actual performance numbers so they can be incorporated into the software simulator.
06/21 Jeff has published numbers for dram load and a dcache fill.

Suspended Items

Item: Unsnap code

Who: lisar, jeffm

Status: Suspended

02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned.
Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump.
We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.
03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.
05/10 Back to life. Does the IKOS support RAM dump?
05/24 This item was suspended again. There are no resources allocated to this item at this time.

Item: When do we have a full calliope simulation available (IKOS)?

Who: lisar

Status: Suspended.

04/26 This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.
05/10 Lisar was not able to attend the meeting.
05/24 Suspended. This is related to the snap/unsnap item (above).

Completed Items

None.

Terp Feature Status

```

inprog   o Ifetch protection granularity
          - performance vrs accuracy tradeoff
inprog   o Fetch instructions as octlets
          o Implement hardware configuration through Cerberus regs
          (SDRAM parameters, dram enable?)
done     o ability for terp to load hermes sections
          - lisar would like this functionality added
inprog   o remove stbio from hwterp.

```

----- line of consciousness -----

```

o Better latency model for Calliope accesses
o Checkpoints/Snapshots

```

Performance Test Status

Status	Test Description	Verified		S/W
		H/W	S/W	
ran	store/load to dram	Yes	No	
ran	store/load to hermes	No	No	
ran	store/load to cerberus	No	No	
ran	load from rom	No	No	
ran	icachemiss	No	No	
ran	dcachemiss	No	No	
	load ltlb entry (write+read)	No	No	
	load gtlb entry (write+read)	No	No	
	NB overflow	No	No	
	generate an interrupt	No	No	
	return from interrupt	No	No	
	mult cyls trying to take exceptions at the same time	No	No	No
	predicted branch	No	No	
	unpredicted branch	No	No	
ran	dcache fill	Yes	No	
	single cylinder exception	No	No	
	cable_in_main_handler-- inner loop in EQ_UPDATE_WEIGHTS() (khp)			No No
	NTSC encode loop (ron)	No	No	
	rs_compute_syndrome() (larry)	No	No	
	IDCT code (gregg)			
	pseudo-Huffman decode loop (gregg)	No	No	
	a reconstruction routine from macroblock.c (gregg)			No No

Classes	Tests Written
Loads	
Stores	s64la
Atomic	saas64la, smas64la
Branches	bi
Gateway (??)	bgate, bgatei
blink{,I} (pred and unpred)	--
{G,E}set	gsete64
{G,E}{add,sub,mux,logical}	eaddi
Easum	easum
Eulms	--
Elms	--
Ggfmul8	ggfmul
G{,U}mul{,add}{8,16}	gmuladd8, gmuladd16
G{,U}mul{,add}32	gmuladd32

G{,U}mul64	--	
G{,U}muladd64		gmuladd64
Gextgract128	--	
Other gextract instr		gextracti128
Expand/Compress	--	
Shift/Rotate	--	
Copy/Swap	--	
Shuffle/Deal	--	
Suffle/Mux		eshufflei4mux
Select	--	
Deposit/Withdraw	--	

Dave Tomlinson has noticed that the stgen (store/load) tests have similar cycle counts between the HW and SW simulators but the regdepend (register dependency) tests take 2 - 3 times as many cycles on the HW as predicted by the SW simulator. Jeff and dit00 are going to analyze a trace.

Test Status and General Discussion

We need to modify the ukernel's taskExit so that tests running on the HW simulators can tell when the test finishes. Lisa continues to look at this.

Hermnasty test found a HW bug.

Hermes interleave tests have some failures. Some in the test and some in the HW

Amber L. Smith

From: lisar (Lisa Robinson)
Sent: Wednesday, June 14, 1995 12:01 PM
To: agc; dbulfer
Cc: tbr
Subject: mnemo stuff

Hi guys

Just to make sure you know that about 3 times a week the my group gets together to go over action items etc with respect to our role in mnemo.

We have (or are just about to have) such a meeting at 10.00am today though usually it is at 11.00am (mon, wed, fri).

You are welcome at any time to stop in, also the open items are tracked on my white board, I can go over these with you at any time.

I understand that we are not making even close to the rate of progress we had intended for this phase of the mnemo and there are many factors involved here but I do want to assure you that the following folks are working towards mnemo tapeout.

veena, deepak, (functional verification) doi (throttle development, pci bridge, convenhancement etc) jeffm (cute dram model) dit00 (re-build of released verification) solo (proteus snapshot maintenace and custom block LVS/DRC) lisar (simulator build, vector generation from hc files, regression, etc)

Lisa R.

Amber L. Allen

From: lisar (Lisa Robinson)
Sent: Wednesday, June 14, 1995 11:46 AM
To: veena; brian; solo; doi; deepak
Cc: geert; tbr; drew; manser
Subject: Meet this morning

Lets get together this morning to discuss the following. The aim is to have prioritized list of tasks covering the next 2 weeks by the end of today.

10.00am Mnemo - veena, deepak, solo, doi, dit00
(veena, deepak please provide a list of the current test status)

10.30am Pollux - veena, brian, solo, dit00

11.00am Cronus - veena, brian, solo, dit00, doi, jeffm

1.00pm Buterpe - dit00, jeffm

For all but the first meeting (which will be as the usual 3 times a week meeting in my office) will meet in the little war room if available or the boxers conference room if not.

Lisa R.

From: tbr
Sent: Wednesday, June 14, 1995 1:20 AM
To: Larry Yamano
Cc: abbott@microunity.com; agc@microunity.com; khp@microunity.com; yam@microunity.com
Subject: MU Chips?

Larry Yamano wrote (on Tue Jun 13):

Hi Curtis, Alan, Tim,

The DSP group (khp in particular) has taken on the task of better defining a CDM design. I believe that only after completing this exercise can a reasonable rough BOM be generated. However, it occurred to me that the "design" MU is "targeting" has shifted again. My current impression (via word of mouth) is that there are two chips:

Thinking changed here for three reasons. First major surgery on Euterpe/cronus to make space for the DSP functionality is a big schedule risk item relative to leaving the current designs alone and defining clio as a separate chip.

Second even if we take the risk it looks like we'd have a hard time taking enough out to accommodate the new functionality without growing the die size. Doing that would probably be less cost effective than adding a second small die (especially in the foundry process).

Third, mouss felt that making a major investment in the single chip solution before the market is established is a business risk since the design is potentially useless for any other application. We assume the work of building clio is much less than the single chip and so if things don't go well we have less to lose and the generic CPUs are potentially re-targettable. However, if things go well and volumes look set to increase we still have the option to go back and integrate and by then we will know for sure just what we can leave out.

Clio:
300MHz
.5u Mobimos
Power with air bridge: 5W
Power w/o air bridge: >3 x 5W = >15W(?)

With the current process non-airbridged power will be higher (maybe as much as 5x) since we will have nitride rather than oxide dielectric.
We have to assume airbridge before we have a real product (at least that's what I think).

Includes:
A/D
Uplink DAC
A/D clock generation
Hermes Channel

Potentially includes the A/D and D/A, however, will also be able to work with external components.

Euterpe:
300MHz
.5u Mobimos
Power with air bridge: 17W

Expect it to be a bit higher than this (25W). We are at 85W for nominal 1GHz however we are likely to tape it out with timing violations remaining, which will require relatively higher power then to make the 300MHz number. We expect to continue to optimize it in later revs (which we expect to have to make as there will be further design rule changes before yield really comes up).

Power w/o air bridge: $>3 \times 17W = >51W(?)$

Expect it to be 85W without airbridge and it may not make 300MHz.

Includes:

Hermes Channel

Schedule:

Possible Euterpe w/o air bridge late 95

Yes.

Probably no working Clio/Euterpe chips for at least 9 months (2nd Qtr 96)

We should not plan for failure here. We have to go by current best estimates.

Air Bridge working ____?

Only al can fill this in.

Some Questions:

1. Are we assuming air bridging works for the first trial units?

Would depend on the scale of the trial I think. We have been discussing re-using the Hestia mechanical design which should be able to handle the power of non airbridged parts.

2. What would the power estimates be for .4u Mobimos?

Assume we go from 25W -> 15W.

3. Could Clio also include the DACs for

a. AGC control?

b. tuner frequency selection?

c. uplink variable amp control?

Yes, but again we'd want the escape of being able to implement these on the outside if the analog portion of the process still has trouble.

4. Could Clio also include the antialias filter prior to the A/D?

One for agc . . .

I limited distribution of this note since many of the above assumptions are probably wrong.

Don't take my answers to be the final word either!

Tim

From: graham (Graham Y. Mostyn)
Sent: Tuesday, June 13, 1995 10:39 PM
To: al; ahn; hopper; andrew; geert; trancy; dane; yves; vil; vanthof
Cc: mudge; manser; mouss; tbr; dbulfer; graham; paulp
Subject: Summary of DAC meeting

Summary of a meeting held on Monday, June 12, to plan the 2-level metal mobimos DAC project.

Attendees:

al, andrew, hopper, geert, vil, dane, yves, vanthof, graham

Project objective:

The primary purpose of the DAC development is to assist fab bring-up with 2 level metal, and provide process characterization and reliability data. Fast tape-out is paramount.

Consequently, the packaged DACs are not designed to emulate any industry standard parts, and may not offer commercial ESD characteristics.

1. Circuit design/layout

In 4 level metal, the DAC core is 0.25 sq mm, with 0.25 sq mm additional for bias. 30 pins for signal and ~14 for multiple supplies suggests a 44 pin package.

Using 2 level metal, a 1mm square circuit area is possible, filled out and pad limited to a 2.2 x 2.2 mm die.

Layout guidelines:

- avoid minimum dimensions when unnecessary
- avoid minimum metal spaces adjacent to large metal areas
- aim for via coverage of >3%
- aim for maximum wiring density over the die
- hold review of individual cells which cover >10% wafer

Action: Dane to arrange circuit design review (requested by Andrew). Ongoing layout reviews to be held regularly over the layout cycle with the complete team.

2. Methodology

Employ Pollux methodology; build baseplate with atoms and drop in custom cells. Consider changing 72 pad (at 150um pitch) Pollux module to 44 pads with >200um pitch.

Al advised using existing frames with proven alignment marks, etc., and a completely filled reticle.

A benign pattern - eg atoms - is used to fill areas not carrying active circuitry, including the 100um scribe.

Circuit designers to better understand post-processing wafflization algorithms, to optimize final manual and automatic waffle results.

Action: Hopper to arrange waffle meeting (!)

3. Packaging

Need to identify contractor and appropriate design rules.

Request packaging options, including standard 44 pin PLCC capability (plastic and ceramic) with gold ball bonding.

MUSE would sort and wafer map.

Action: Graham to contact Indy in Manteca; also Spectrum locally.

4. Fabrication

via 2-3 (220 mask) becomes a passivation (pad) mask.

Action: Need simple amendment to current design rules to accommodate process changes (A1).

5. Testing

Action: Test engineering to obtain probe cards; and have test pattern available concurrently with circuit layout (Andrew).

[Note: need to develop burn-in capability, not discussed @ meeting].

6. Schedule

Assumptions: * 2 mask designers allocated to project.
 * fracturing time of order 2 days only

All times are cumulative from project commencement. Aggressive timing.

Circuit design and layout: t=4 weeks

Test patterns ready: t=4weeks

CAD (interstitial fill, DRVC, LVS clean): t=6 weeks Tape out: t=6 weeks First mask received: t=7 weeks Start process: t=8 weeks Processing time gated by masks received; goal = 1 mask/day.

Wafers out: t=11 weeks

Proceed to wafer sort; package; burn-in.

There was discussion of eliminating bias circuit and reducing layout by 2 weeks. Design team felt that the loss of data and more complex necessary external bias set-up was a poor trade for the shortened design cycle.

7. Resources

The allocation of 2 full time mask designers for 4 weeks would significantly impact existing tape outs. All other disciplines would experience finite but perhaps lesser impact to existing projects.

Action: Graham/Geert to pursue priority determination with management.

From: tbr
Sent: Tuesday, June 13, 1995 12:02 PM
To: brianl; geert; hopper; lisar; tom; wingard
Subject: atlas cell modelling

At lisar's request, I'd like to reschedule this to 1pm. She has a conflict with an ergonomics class. Sorry for the short (ie negative) notice.

Tim

tbr wrote (on Mon Jun 12):

As I understand it there is a possibility that the atlas leaf cells are constructed in such a way that we may run into the same problem attempting to simulate the Cronus LVS netlist as we have done with certain custom leaf cells in Euterpe. The problem relates to the way internal nodes get tied off to supply rails, and the way the spice to edif conversion handles these. Can we have a brief meeting Tuesday 10am please to check if this is a real problem, and if so how we are going to handle it.

Thanks
Tim

Subject: atlas

From: tom (Tom Laidig [tau])
Sent: Tuesday, June 13, 1995 12:00 PM
To: Tim B. Robinson
Cc: lisar (Lisa Robinson); brianl (Brian Lee); wingard (Drew Wingard); geert (Geert Rosseel);
hopper (Mark Hofmann); tau
Subject: Re: atlas cell modelling

Tim B. Robinson writes:

As I understand it there is a possibility that the atlas leaf cells are constructed in such a way that we may run into the same problem attempting to simulate the Cronus LVS netlist as we have done with certain custom leaf cells in Euterpe. The problem relates to the way internal nodes get tied off to supply rails, and the way the spice to edif conversion handles these. Can we have a brief meeting Tuesday 10am please to check if this is a real problem, and if so how we are going to handle it.

I'm already scheduled for my ergonomics class today at 10.

--

τ

From: tbr
Sent: Tuesday, June 13, 1995 1:00 AM
To: stick; vanthof
Subject: forwarded message from tbr

Sorry, missed copying you on this . . .

----- Start of forwarded message -----

Status: RO

X-VM-v5-Data: ([nil nil nil nil nil nil nil nil nil]

["514" "Mon" "12" "June" "95" "17:53:55" "PDT" "tbr" "tbr" nil "13" "atlas cell
modelling" "^To:" nil nil "6"])

To: lisar,brianl,wingard,tom

cc: geert,hopper

Subject: atlas cell modelling

As I understand it there is a possibility that the atlas leaf cells are constructed in such a way that we may run into the same problem attempting to simulate the Cronus LVS netlist as we have done with certain custom leaf cells in Euterpe. The problem relates to the way internal nodes get tied off to supply rails, and the way the spice to edif conversion handles these. Can we have a brief meeting Tuesday 10am please to check if this is a real problem, and if so how we are going to handle it.

Thanks

Tim

----- End of forwarded message -----

From: doi (Derek Iverson)
Sent: Monday, June 12, 1995 8:22 PM
To: guarino; gmo; doi; gregg; claseman; lisa; jeffm; ethan
Cc: hestia
Subject: Software Bringup Meeting Minutes - June 7, 1995

Software Bringup Meeting

June 7, 1995

Next Meeting: June 14 at 2:00 pm.

Attendees: guarino, gmo, doi, gregg, claseman, lisa, jeffm, ethan

New Action Items

Item: Build osf kernel and perpare it for execution on HW simulator
Who: doi/iimura
Status: New

Review of Action Items

Item: Plan for testing remote debugging environment
Who: everyone
Status: Pending

04/13 This will be discussed at the next meeting.
04/20 Short discussion on the pieces required and if we want to have
a standalone version of the hostio software.
Wally has started on this already.
04/26 Any work with snoopy/dram models will begin after euterpe tapeout
when jeff has more time.
05/03 Wayne F. expects to have a prototype ISA card ready for testing
tomorrow. Gmo is in the process of writing code to test the card.
05/10 Gmo has written a test program. Gmo and Wayne need to get together
and try it on the board.
05/24 Wally has added more connectivity between the ukernel and the
debugger (ability to interrupt kernel via the cerberus
'forced interrupt' bit).
Gmo is ready with a test when Wayne has the board ready.
05/31 Wayne has card ready to test (fpga that talks to ISA) for initial
testing. Documentation for the interface next week some time.

Item: ukernel needs to detect machine checks and 'do the right thing'
Who: lisa
Status: Pending

04/13 No new progress.
04/20 On list of things to do. Lower priority.

Item: Create performance test plan
Who: claseman
Status: [11/30] No progress as focus is on functionality.

11/30 We continue to run tests to help us compare terp vs hardware

performance.

We still need to put together the actual performance tests that need to be run on the hardware.

- 05/10 Back to life. Tim has checked in more tests. We are going to use the tests written for hermes and cerberus read accesses in the 'How to debug Euterpe' presentation on monday. We will be looking for the time it takes from the instruction issue, entry to nb, and completion of access.
- 05/24 Tim is going to send lisar a list of the test names that he intends to write so they can be incorporated into the test template.
- 05/31 Jeffm is going to analyze approx one test per day (with Tim looking over his shoulder) to determine/verify the actual performance numbers so they can be incorporated into the software simulator.

Suspended Items

Item: Unsnap code
Who: lisar, jeffm
Status: Suspended

- 02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned. Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump. We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.
- 03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.
- 05/10 Back to life. Does the IKOS support RAM dump?
- 05/24 This item was suspended again. There are no resources allocated to this item at this time.

Item: When do we have a full calliope simulation available (IKOS)?
Who: lisar
Status: Suspended.

- 04/26 This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.
- 05/10 Lisar was not able to attend the meeting.
- 05/24 Suspended. This is related to the snap/unsnap item (above).

Completed Items

Item: Need verify/ukernel Makefile modified to build other tests
Who: doi
Status: Done

Item: Modify startup code in stb/stand to read the cerberus node number
Who: gmo
Status: Done

- 05/03 No new progress

Terp Feature Status

```

inprog      o Ifetch protection granularity
              - performance vrs accuracy tradeoff
inprog      o Fetch instructions as octlets
inprog      o Accuracy wrt HW simulator(s?)
              o Better latency model for Calliope accesses
              o Implement hardware configuration through Cerberus regs
                (SDRAM parameters, dram enable?)
              o Checkpoints/Snapshots
inprog      o ability for terp to load hermes sections
              - lisar would like this functionality added
inprog      o remove stbio from hwterp.
  
```

Performance Test Status

Status	Test Description	Verified			
		H/W	S/W		
ran	store/load to dram	No	No		
ran	store/load to hermes	No	No		
ran	store/load to cerberus	No	No		
ran	load from rom	No	No		
ran	icachemiss	No	No		
ran	dcachemiss	No	No		
	load ltlb entry (write+read)	No	No		
	load gtlb entry (write+read)	No	No		
	NB overflow	No	No		
	generate an interrupt	No	No		
	return from interrupt	No	No		
	mult cyls trying to take exceptions at the same time	No	No	No	No
	predicted branch	No	No		
	unpredicted branch	No	No		
inprog	One instruction from each instruction class:			No	No
	arith store eshifty imul64				
	branch swap epop idiv				
	branchx stored imul4 gfmul				
	load eset imul8				
	loadx eshift imul16 atomic_ops				
	nbload eshiftx imul32				
ran	Instructions that have tests written already			No	No
	eaddi bi bgatei s64la				
	saas64la gsete64 eshufflei4mux gestracti64				
	gesxtracti128 easum gmuladd8 gmuladd16				
	gmuladd32 gmuladd64 ggfmul bgate				
	smaas64la				
inprog	Inner loop sequences:			No	No
	- cable_in_main_handler-- inner loop in				
	EQ_UPDATE_WEIGHTS() (khp)				
	- NTSC encode loop (ron)				
	- rs_compute_syndrome() (larry)				
	- IDCT code (gregg)				
	- pseudo-Huffman decode loop (gregg)				
	- a reconstruction routine from macroblock.c (gregg)				

Test Status and General Discussion

We need to modify the ukernel's taskExit so that tests running on the HW simulators can

tell when the test finishes. Lisa is looking at this.

The bug arrival rate is still pretty constant (about 1 per day).

Hermes interleave is working properly now.

There are still some issues related to the cache controller.

The use of sync ops in the nasty tests are finding problems.

Two dram timing ratios are not working as advertised.

From: tbr
Sent: Monday, June 12, 1995 7:54 PM
To: lisar; brianl; wingard; tom
Cc: geert; hopper
Subject: atlas cell modelling

As I understand it there is a possibility that the atlas leaf cells are constructed in such a way that we may run into the same problem attempting to simulate the Cronus LVS netlist as we have done with certain custom leaf cells in Euterpe. The problem relates to the way internal nodes get tied off to supply rails, and the way the spice to edif conversion handles these. Can we have a brief meeting Tuesday 10am please to check if this is a real problem, and if so how we are going to handle it.

Thanks
Tim

From: hopper (Mark Hofmann)
Sent: Friday, June 09, 1995 3:49 AM
To: vant
Cc: geert (Geert Rosseel); tbr (Tim B. Robinson); vanthof (Dave Van't Hof); tom (Tom Laidig)
Subject: Re: minimum diffusion xtors in euterpe

vant writes:

Last night I ran a test on euterpe which searched for all minimum devices and flagged diffusion areas which would have to be added to meet the new proposed rules from Paul and Al. I did this in two steps; the first flagged all devices where width = .5microns, the second flagged all devices from .55 to .75 microns. The results are quite promising. I'd say that over 95% of the min devices (.5microns) are in memory blocks (cache, ctag, gtlb, and cr) while the rest are mostly in leaf cells. The .55 - .75 micron devices all appear to be in memory blocks.

All of this is based on the assumption that fixing the minimum devices does not cause other drc errors. Since most of the flags were in the cache and Stick already has a local version with the correct fixes for the memory cell, we may get lucky and not have any major layout disasters trying to fix the remaining errors.

Dave

Okay. Thanks for the run Dave.

Of the 5% of min devices not in the cache, do you have a rough idea how many cells are effected? How many devices are we talking about?

-thanks,
hopper

Antony, Mike

From: vanthof (vant)
Sent: Friday, June 09, 1995 9:01 AM
To: hopper (Mark Hofmann); geert (Geert Rosseel); tbr (Tim B. Robinson)
Cc: vanthof (Dave Van't Hof); tom (Tom Laidig)
Subject: minimum diffusion xtors in euterpe

Last night I ran a test on euterpe which searched for all minimum devices and flagged diffusion areas which would have to be added to meet the new proposed rules from Paul and Al. I did this in two steps; the first flagged all devices where width = .5microns, the second flagged all devices from .55 to .75 microns. The results are quite promising. I'd say that over 95% of the min devices (.5microns) are in memory blocks (cache, ctag, gtlb, and cr) while the rest are mostly in leaf cells. The .55 - .75 micron devices all appear to be in memory blocks.

All of this is based on the assumption that fixing the minimum devices does not cause other drc errors. Since most of the flags were in the cache and Stick already has a local version with the correct fixes for the memory cell, we may get lucky and not have any major layout disasters trying to fix the remaining errors.

Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: tbr
Sent: Thursday, June 08, 1995 12:57 AM
To: geert (Geert Rosseel)
Subject: euterpe

Geert Rosseel wrote (on Wed Jun 7):

Hi Tim,

I am trying to find out what my problems are and I think I may have a bad chip_euterpe-base.strength file.

Can you check the one that is checked in and see how good it is ..

The one checked in is bogus because there have been changes to the sub blocks. I was assuming that since you said you were rebuilding the subblocks you would touch rebuild and have it regenerate these files from scratch. (That's what I did). I have updated to BOM 318 to get today's bug fix. What I have not seen is a placement change from woody that I was expecting to fix the one placement problem in hc. I will shortly start up my run and if it does get through placement I will end up being able to check in new base.strength etc.

P.S. There is another meeting tomorrow at 4:00 about the same subject.

Sorry I had to leave. How long did it go on?

The current thinking seems to be to have a dual tract

Have a system designed on euterpe and clio

Both chips designed on mobi and outside

Which comes first?

Use the Hestia system as a base for the box design

Clio is a digital only chip (200.000 atoms in mobi)

Have prototypes based on
Mobi Euterpe and Clio
Or

Foundry Euterpe and Clio

Power-reduce Mobi design with airbridge and re-optimize for
300MHz design

Power-reduce Foundry designs with new CMOS micro-architecture.

Tomorrow there is going to be discussion about resources
and priorities *which designs first ?)

I guess that answers my question above!

Tim

From: geert (Geert Rosseel)
Sent: Wednesday, June 07, 1995 10:22 PM
To: tbr
Subject: euterpe

Hi Tim,

I am trying to find out what my problems are and I think I may have a bad chip_euterpe-base.strength file.

Can you check the one that is checked in and see how good it is ..

P.S. There is another meeting tomorrow at 4:00 about the same subject.

The current thinking seems to be to have a dual tract

Have a system designed on euterpe and clio

Both chips designed on mobi and outside

Use the Hestia system as a base for the box design

Clio is a digital only chip (200.000 atoms in mobi)

Have prototypes based on
Mobi Euterpe and Clio
Or
Foundry Euterpe and Clio

Power-reduce Mobi design with airbridge and re-optimize for
300MHz design

Power-reduce Foundry designs with new CMOS micro-architecture.

Tomorrow there is going to be discussion about resources and priorities *which designs first ?)

Geert

From: hopper (Mark Hofmann)
Sent: Wednesday, June 07, 1995 1:37 AM
To: tbr (Tim B. Robinson); geert (Geert Rosseel)
Subject: 4 tapes sent to DuPont today (fwd) -FYI

hi,

Just to let you know what's going out :

Kurt Wampler writes:

To: al, fung, paulp
Subject: 4 tapes sent to DuPont today
Cc: hopper, wampler

I sent all 4 layers to DuPont this afternoon:

Orchis 060B (0009-XX-060B1-1-0150)
150C (0009-01-150C1-1-0151)

Calliope1 060B (0004-XX-060B1-1-0152)
150C (0004-01-150C1-1-0153)

I alerted Bob Hickman to expect the larger serifs on both masks, and also that there are known 0.25uM slivers on the 150C masks that may potentially show up as false defects during inspection.

The patterns for these masks have been installed in the online pattern archives on CimPhoto5.

- Kurt

From: tbr
Sent: Tuesday, June 06, 1995 5:07 PM
To: woody (Jay Tomlinson)
Cc: geert (Geert Rosseel); vo (Tom Vo)
Subject: Re: Possible wire savings

Jay Tomlinson wrote (on Tue Jun 6):

Tom Vo wrote (on Tue Jun 6):

Tim B. Robinson wrote

>

>

>With the new ecl2cmos cell, topt is reporting

>

>Warning! Instance iol/Infifo/Urawd1/u0 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd1/u1 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd1/u2 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd1/u3 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd1/u4 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd1/u5 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd1/u6 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd1/u7 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u0 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u1 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u2 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u3 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u4 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u5 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u6 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>Warning! Instance iol/Infifo/Urawd0/u7 (xbbufdf2s) could be HALF swing, but is being forced to FULL swing.

>

>and similarly for io0. I'm not sure if these wires go directly into
>the Cerberus choke point, or if we have already muxed them on the
>outside. At any rate, if we think reducing the wire counts in these

Those wires are already muxed on the outside . Converting them would help the wiring in the io0,io1 area though .

It certainly can't hurt the routing around hc0/nb and uu.
woody

Then I suggest we try it. I'm game in my local copy. However, will we need an explicit vref generator since the converter is in cmos sofa and may not be able to be hooked up in the normal way if we just leave te input floating?

Tim

Exhibit C15

From: stick (Bruce Bateman)
Sent: Tuesday, June 06, 1995 12:02 PM
To: geert
Subject: Re: Layout reviews

> Date: Tue, 6 Jun 1995 08:11:48 -0700
> From: geert (Geert Rosseel)
> To: ahn, al, bill, bpw, geert, hopper, ong, paulp, stick, tau, vanthof, vo,
> wampler, wingard
> Subject: Layout reviews
> Cc: mouss
>
> Hi,
>
> We'll have some reviews of Euterpe related layouts by the fab-group.
>
> We'll start with the euterpe memory cell structures. The first review
> will be tomorrow (wednesday) at 4:45 in the War Room (after Al's group meeting).
>
> Bruce : can you make a set of plots of the memory cell for the meeting.
>

No problem.

BB

Archival Label

From: geert (Geert Rosseel)
Sent: Tuesday, June 06, 1995 10:12 AM
To: ahn; al; bill; bpw; geert; hopper; ong; pulp; stick; tau; vanthof; vo; wampler; wingard
Cc: mouss
Subject: Layout reviews

Hi,

We'll have some reviews of Euterpe related layouts by the fab-group.

We'll start with the euterpe memory cell structures. The first review will be tomorrow (wednesday) at 4:45 in the War Room (after Al's group meeting).

Bruce : can you make a set of plots of the memory cell for the meeting.

Geert

From: efelias (Eldred Felias)
Sent: Tuesday, June 06, 1995 1:35 AM
To: bpw (B. P. Wong); geert (Geert Rosseel)
Subject: Re: please unlock the following cells... (fwd)

Never mind about the previous message. Dave just answered my question.

Eldred

Forwarded message:

>From: vanthof Mon Jun 5 23:31:10 1995
>From: vanthof (vant)
>Message-Id: <199506060631.XAA00799@hestia.microunity.com>
>Subject: Re: please unlock the following cells...
>To: efelias@hestia.microunity.com (Eldred Felias)
>Date: Mon, 5 Jun 95 23:31:08 PDT
>Cc: vanthof (Dave Van't Hof)
>In-Reply-To: <199506060626.XAA12206@poseidon.microunity.com>; from
>"Eldred Felias" at Jun 5, 95 11:26 pm
>X-Mailer: ELM [version 2.3 PL11]
>
>Eldred Felias writes:
>>
>>...gtxrdsa, gtxordiff, gtmpl3. All of these cells have the m1 concave
>>minimum violation of 30 udrs. I also have a bunch on the edges where
>>m1 hooks up to the m1 of narborderpwr and I don't know (yet) how to
>>fix them. You know, I am sure the gt1b was clean already i.e., it didn't contain any of
>>these "concave"
>>violations. Is this check fairly new?
>>
>>
>>Eldred
>
>Yes, this check was recently extended from 20 to 30 udrs and lots of
>things are being flagged. We have a way to fix these at the top level
>at tapeout, so fix the easy ones and punt on the difficult ones, but do
>get as many as possible so the top level work is minimal.
>
>The cells are unlocked.
>
>Thanks,
>Dave
>
>---
>Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
>vanthof@microunity.com 1 408 734-8100
>"Don't blame me! I didn't vote for him"
>

--

From: tbr
Sent: Saturday, June 03, 1995 4:28 PM
To: doi
Cc: jeffm; tom
Subject: Urgent nit problem

A rebuild of the euterpe tapeout snapshot proteus has died because of being unable to regenerate in proteus/verilog/zlibsrc. It appears the root cause of the problem is the following line in the Makefile which violates the absolute rule of no absolute paths pointing outside the BOM:

```
TERPSWDIR = /p/soft/stb/sun4-src/gnu-tools/sim/terp
```

Something out of the control of the BOM has changed and logicsim will no longer compile. Please clean this up and avoid the temptation to use this sort of kludge in the future. We're dead in the water till we can get past this.

Tom, can you run another nit-hunt over euterpe and proteus and see if there is anything else like this waiting to trip us up.

Tim

From: wingard (Drew Wingard)
Sent: Friday, June 02, 1995 7:13 PM
To: geert
Subject: Re: Back from meeting...

> We concluded that we're a week late on about everything (except
> Brain's stuff). So we moved the tentative tape-out date with a week to
> July 22. I still need to write up a summary.

>
> Geert
>

Great! Maybe I can get a little sleep this weekend...

I heard something scary over lunch: IBM and other C4 folks are having significant worries over alpha particles from the lead solder bumps. Apparently IBM is concerned enough that it went out and purchased a lead mine that happens to have the world's lowest concentrations of the radioisotope that causes the emissions.

And the jury is still out on on-chip decoupling caps with respect to yield. But the latest DEC Alpha design adds ~ 160nF using gate oxide under the supply bussing...

Drew

From: solo (John Campbell)
Sent: Friday, June 02, 1995 1:02 PM
To: solo
Cc: yves (Jean-Yves Michel); vanthof (Dave Van't Hof); tau; tbr (Tim B. Robinson); geert (Geert Rosseel); graham (Graham Y. Mostyn); lisar (Lisa Robinson)
Subject: Re: p4s.ly

as solo was saying

..
..Looks like you made an unreleased change to p4s. you should be aware ..that this layout is used extensively in chips near tapeout. we may ..want to consider renaming the cell for your purposes.

..
..let me know what was done to this cell and please don't release it ..without the approval of tbr or myself.

.....
..regards,
..solo a.k.a. John Campbell x516

..
..
..
..
these cells are ok for drc on this change. some will need to be changed because of the the new notch flow. i can go with releasing the p4s.ly if tim wants to include it in the next snapshot.

....
regards,
solo a.k.a. John Campbell x516

From: yves (Jean-Yves Michel)
Sent: Friday, June 02, 1995 11:26 AM
To: solo
Cc: vanthof; tau; tbr; geert; graham; lisar; tom
Subject: Re: p4s.ly

> From solo Fri Jun 2 08:27:56 1995
> From: solo (John Campbell)
> Subject: p4s.ly
> To: yves (Jean-Yves Michel)
> Date: Fri, 2 Jun 95 8:27:53 PDT
> Cc: vanthof (Dave Van't Hof), tau, tbr (Tim B. Robinson),
> geert (Geert Rosseel), graham (Graham Y. Mostyn),
> lisar (Lisa Robinson)
> X-Mailer: ELM [version 2.3 PL11]
> Content-Length: 349
>
> Looks like you made an unreleased change to p4s. you should be aware
> that this layout is used extensively in chips near tapeout. we may
> want to consider renaming the cell for your purposes.
>
> let me know what was done to this cell and please don't release it
> without the approval of tbr or myself.
>
> regards,
> solo a.k.a. John Campbell x516
>
>
>
>

I did a grep across proteus to find out who is using it and I've warned the other designers of this change. Is this transistor used in another database than proteus?

As for the change, when this transistor was array'd in a certain pattern, it was showing a col. plug spacing violation between 2 instances placed side-by-side and mirrored. This is against our rules for building the xtor library. I noticed that the collector plug geometry was 10 x 12 udr where other PMOS xtor had a 10 x 10 plug. So I cut the plug to 10 x 10 without changing anything else. I don't think anybody will be affected by that.

If it is a problem to release it with euterpe or mnemo (if p4s is used there), then I can create a new xtor to fix my pollux layout.

Let me know what is the best.

Jean-Yves

Archival. Add -

From: tom (Tom Laidig [tau])
Sent: Friday, June 02, 1995 11:21 AM
To: John Campbell
Cc: yves (Jean-Yves Michel); vanthof (Dave Van't Hof); tau; tbr (Tim B. Robinson); geert (Geert Rosseel); graham (Graham Y. Mostyn); lisar (Lisa Robinson)
Subject: Re: p4s.ly

John Campbell writes:

Looks like you made an unreleased change to p4s. you should be aware that this layout is used extensively in chips near tapeout. we may want to consider renaming the cell for your purposes.

let me know what was done to this cell and please don't release it without the approval of tbr or myself.

Looking at the files, the change is just a reduction in the size of a well tie from 10x12 udr to 10x10 udr. I don't know what motivated the change (perhaps the previous well tie was close enough to the cell edge to have some interaction with another cell?) but it looks quite benign as far as I can tell.

--

'\u'

Andrew L. Alden

From: solo (John Campbell)
Sent: Friday, June 02, 1995 10:28 AM
To: yves (Jean-Yves Michel)
Cc: vanthof (Dave Van't Hof); tau; tbr (Tim B. Robinson); geert (Geert Rosseel); graham (Graham Y. Mostyn); lisar (Lisa Robinson)
Subject: p4s.ly

Looks like you made an unreleased change to p4s. you should be aware that this layout is used extensively in chips near tapeout. we may want to consider renaming the cell for your purposes.

let me know what was done to this cell and please don't release it without the approval of tbr or myself.

....
regards,
solo a.k.a. John Campbell x516

From: hopper (Mark Hofmann)
Sent: Thursday, June 01, 1995 6:02 AM
To: Tim B. Robinson
Cc: vanthof (Dave Van't Hof); hopper@microunity.com
Subject: Re: euterpe drc run

Tim B. Robinson writes:

Sorry, you should have been in the loop on this and I think I only copied solo figuring his automated stuff would be the first to meet this. As I understand it it's a "trivial" metal edit, to bring an internal VREF line to a target so we can drive it from the outside. It will be done today, and we will update the snapshot proteus and the euterpe baseplate snapshot as soon as both schematic and layout change are completed.

Okay, no problem.

I think, as Dave points out, another VREF generator will be added to the Sofa as well to power the gtlb. But still, it should be a metla only change.

-hopper

Author: [REDACTED]

From: tbr
Sent: Thursday, June 01, 1995 12:56 PM
To: vanthof (vant)
Cc: Mark Hofmann; vanthof (Dave Van't Hof)
Subject: Re: euterpe drc run

vant wrote (on Thu Jun 1):

Mark Hofmann writes:

>
>vant writes:
> Tim,
> I've just installed the new drc flow, however, there was one other
> cell which is still being edited. Once that is done, then I can start
> up a round of euterpe drc's. The cell edits will be done tomorrow.
> There are also many pad cell edits going on. Those will take about a
> week, but I'm not going to wait for those to be completed. besides,
> I can get one complete drc run in before the pad edits are completed.
>
> sounds like a good idea, dave.
>
> -hopper
>

Turns out there may be more edits. Jay just asked me about some edits to the gtlb and when they were going to happen. I know nothing about them but it sounds like it's layout edits, which pretty much invalidates any attempts at running drc's until the baseplate is rebuilt.

Sorry, you should have been in the loop on this and I think I only copied solo figuring his automated stuff would be the first to meet this. As I understand it it's a "trivial" metal edit, to bring an internal VREF line to a target so we can drive it from the outside.

It will be done today, and we will update the snapshot proteus and the euterpe baseplate snapshot as soon as both schematic and layout change are completed.

Tim